SCO® UNIX® System V/386

Operating System

Upgrade Notes

Release 3.2 Version 2.0

The Santa Cruz Operation, Inc.

Document Version: 3.2v2.0C
Date: 17 May 1990

SCO UNIX System V/386 Release 3.2 Version 2.0 Upgrade Notes

SCO UNIX System V/386
Operating System
Upgrade Notes
Release 3.2 Version 2.0
17 May 1990

1. Overview

These notes describe the procedure for upgrading your current SCO UNIX System V/386 Operating System to Release 3.2 Version 2.0. It also describes the procedure for restoring your current system environment files to save time when reconfiguring your upgraded system.

The upgrade process consists of the following stages:

- Prepare for the upgrade, which includes making a full backup of your system and, if necessary, recording your filesystem information.
- Perform the upgrade.
- Relink the kernel, if necessary.
- Restore specific configuration files.

Note

As part of the upgrade process, you have to reinstall any applications and device drivers you had installed. Be sure and have these on hand before starting the upgrade.

1.1 About These Notes

Please read through the entire document before upgrading the operating system.

SCO is always pleased to hear of users' experiences with our products and to receive feedback on how they can be made even more useful. All written suggestions are given serious consideration.

2. Preparing for the Upgrade

The following procedure explains how to prepare your system for the upgrade process.

Note

This upgrade does not affect any other operating systems (or partitions), such as MS-DOS, that currently share the hard disk, unless you choose to change partition sizes. Refer to the "Using DOS and OS/2" chapter in the *System Administrator's Guide* if you want to change the size of your UNIX partition.

2.1 Shut Down the System

You should be the only person logged in to your computer when you perform the upgrade. Make sure users know what you are going to do, and when you are going to do it. They may wish to make their own backups before you bring down the system.

Also, advise users to remove unnecessary files. This makes the whole process faster and requires fewer floppies (if they are your only backup media).

To shut the system down and enter single user (maintenance) mode, log in as *root* and enter the following command:

/etc/shutdown

Give users five minutes to log off. When you are prompted, reboot the system and enter single-user mode.

2.2 Back Up the System

Before you do anything else, make a full backup of the system. It is recommended that you use the Backups→Create→Unscheduled selection of the **sysadmsh**(ADM). This is the best way to create backups.

You should make separate backups of all non-root filesystems. (*/u* is an example of a non-root filesystem.)

You should make a backup immediately before updating so that you have the latest possible version of your system. This is especially important in the case of user files, which can change frequently. Once again, remove any old, unnecessary files. This makes the whole process faster.

You can use any backup media, including floppy and tape drives, to save your files. However, if you use a special driver with your backup device, you cannot use that device immediately after installing the new kernel, because it is not yet linked with the required device driver. This is important because you can only read information from the default floppy drive when you relink the kernel.

Make certain you have the original distributions of any applications packages you use. You must reinstall those software packages after updating your UNIX System.

2.3 If You Plan to Preserve Non-Root Filesystems

You can leave any special filesystems on your primary and secondary hard disks untouched during the upgrade. You must follow certain precautions, however. When the new root filesystem is created, all the device nodes and the mount points are destroyed, but the actual files and divisions are still intact. (The device node is the device name that refers to the filesystem, such as */dev/u*. The mount point is the directory that the device is "mounted" upon, such as */u*.)

During the upgrade, you are dropped into the **divvy** menu and asked to name (using the "n" option) each of your filesystems over again. (Later you will also run **mkdev fs** again to recreate the device nodes and mount points. You are cautioned to take care not to obliterate or change the filesystems in any way other than to give them names once again. The **divvy** menu has an "r" (restore) option to undo any changes, but you are warned to take care during this part of the procedure.

If you are increasing your swap space or the size of your bad track table, you need to recreate your non-root filesystems. If you are preserving any filesystems, you need to rename your non-root filesystems. To do either, you must record your present **divvy** table. Enter the following command:

/etc/divvy -c 1 -b 1 -p 0

Copy the table entries displayed on your screen onto the blank Table 1 (you need not copy the New File System column shown on your screen).

Table 1: Primary Hard Disk Filesystems

Name	Type	New FS	#	First Block	Last Block

Enter **q** to quit out of **divvy**, followed by **e** when asked if you want to exit.

If you have a secondary hard disk with existing UNIX filesystems, they are not touched by the upgrade, but you must record the filesystems in your present **divvy** table and later rename the divisions after the upgrade is complete. Enter the following command:

/etc/divvy -c 1 -b 1 -p 1

Copy the table entries displayed on your screen onto the blank form in Table 2 (you need not copy the New File System column shown on your screen).

Table 2: Secondary Hard Disk Filesystems

Name	Type	New FS	#	First Block	Last Block

Enter **q** to quit out of **divvy**, followed by **e** when asked if you want to exit.

3. Upgrade Procedure

This section describes how to upgrade your system.

1. You should have your complete UNIX distribution, applications, device drivers, and so forth on hand before beginning. Make certain you have your serialization codes and activation keys as well.

2. You should have performed a complete safety backup of your system as instructed in "Back Up the System."

3. If you have additional filesystems, you should have recorded your **divvy** tables as described earlier in "If You Plan to Preserve Non-root Filesystems."

4. After you notify users that your system will be down for some time, log in as *root* and use the **shutdown**(ADM) command to halt the system:

   /etc/shutdown

5. When the system halts, insert the BOOT floppy (N1) into the drive and press any key to reboot the system. The

system boots from the floppy and the usual messages are displayed. Insert volumes as instructed and select your keyboard nationality as usual. Select **1**, the fully configurable Hard Disk Initialization option.

6. After the hard disk initialization message is displayed, you see the following:

```
During installation you may choose to overwrite all
or part of the present contents of your hard disk.
Do you wish to continue? (y/n)
```

You must respond **y**.

7. Unless you have installed a new hard disk, when the "Hard Disk Drive Configuration" menu appears, simply quit out of it.

8. Next, the **fdisk** menu is displayed:

```
1. Display Partition Table
2. Use Entire Disk for UNIX
3. Use Rest of Disk for UNIX
4. Create UNIX Partition
5. Activate Partition
6. Delete Partition
Enter your choice or 'q' to quit:
```

If you are preserving filesystems on the primary hard disk, then you must quit out of this menu.

If you have a DOS partition on your hard disk, enter **3** and give the starting and ending blocks of the UNIX partition (be sure not to overlap your DOS partition). Otherwise, enter **2** to use the entire hard disk for the UNIX partition. You have now set up the partition(s) to use on your hard disk. To continue with the next step in the procedure, enter **q** and press ⟨Return⟩.

Note

If you have a SCSI or SMS-OMTI controller, you should skip steps 9-18. **badtrk** is not run with disks attached to these controllers.

---

9. Next, this message is displayed if you are preserving filesystems:

```
This device contains a valid division table.  Additional
(non-root) filesystems can be preserved across this reinstallation.
If you wish to be able to preserve these file systems later, you must
not change the current limit of the bad track table, which is
n bad tracks.  Do you wish to leave it unchanged? <y/n>:
```

You must respond **y**. This ensures that the limit of your bad track table is not increased, which would destroy your filesystem(s). You can still add entries to your current bad track table, but you are not allowed to increase the size of the table itself.

The main **badtrk** menu is displayed:

```
1. Print Current Bad Track Table
2. Scan Disk (You may choose Read-Only or Destructive later)
3. Add Entries to Current Bad Track Table by Cylinder/Head Number
4. Add Entries to Current Bad Track Table by Sector Number
5. Delete Entries Individually from Current Bad Track Table
6. Delete All Entries from Bad Track Table

Please enter your choice or 'q' to quit:
```

Enter **2**, then press ⟨Return⟩

10. You see the following submenu:

```
1.    Scan entire UNIX partition
2.    Scan a specified range of tracks
3.    Scan a specified filesystem
```

Select option **1** to scan the whole partition. If you choose option 2 or 3, you are prompted to specify the area you want **badtrk** to scan.

11. After you select the area you want scanned, you are given the choice:

```
1. Quick scan (approximately 7 megabytes/min)
2. Thorough scan (approximately 1 megabyte/min)
```

Select option **2**.

12. You are prompted:

```
Do you want this to be a destructive scan? (y/n)
```

Unless you have installed a new hard disk, you must enter **n**. The non-destructive scan reads but does not write to the disk. If you enter **y**, you are warned:

```
This will destroy the present contents of the
region you are scanning.

Do you wish to continue? (y/n)
```

When scanning begins, you see this message:

```
Scanning in progress, press 'q' to interrupt
at any time.
```

13. After you have responded to the above prompts, the program scans the active partition of the new disk for flaws. The larger your disk, the longer the scanning process takes, so a very large disk may take a while.

As **badtrk** scans the disk, it displays the number of each track it examines, and the percentage of the disk already scanned. Pressing the q key at any time interrupts the scan. If you press q to interrupt the scan, you do not need to press ⟨Return⟩. You are then prompted to continue scanning or to return to the main menu.

Whenever **badtrk** finds a defective track, it lists the location of that track using both the sector number and cylinder/head conventions. Defective track information is entered into the table and displayed on the screen. An example bad track might be:

```
wd: ERROR : on fixed disk ctlr=0 dev=0/47 block=31434 cmd=00000020
    status=00005180, sector = 62899, cylinder/head = 483/4
```

14. When the scan is complete, the menu reappears. Select option **1** to see the results of the scan. Your bad track table might look like this:

```
Defective Tracks
+-----------------------------------------+
|    Cylinder   Head   Sector Number(s)   |
+-----------------------------------------+
|1.     190      3       12971-12987      |
+-----------------------------------------+
Press <RETURN> to continue
```

Press ⟨Return⟩ to return to the main menu.

Note

If there is a flaw in the first few tracks of the UNIX partition, you are returned to the **fdisk** utility (see the previous installation step). Repartition the disk with **fdisk** so that the UNIX partition no longer includes the defective tracks. You will have to experiment to determine how many tracks to exclude. Leave these defective tracks unassigned to any operating system. When you leave **fdisk**, **badtrk** runs again. Scan the disk for flaws. This process continues until **badtrk** finds no flaws in the first few tracks.

15. When you are finished making changes to the bad track table, enter **q** and press ⟨Return⟩ to return to the main menu.

    At the main **badtrk** menu, enter **q** again and press ⟨Return⟩.

16. If you specified at step 9 that you did not want to change the current limit of the bad track table, then you are not prompted to enter a new limit; proceed to step 18. Otherwise, you are next prompted for the number of tracks to allocate as replacements for those tracks that are flawed. You should allocate at least as many as the recommended number. Enter the number or just press ⟨Return⟩ to use the recommended number that is displayed:

```
Enter the number of bad tracks to allocate space for
(or press return to use the recommended value of n):
```

    This number is based on the number of bad tracks currently in the table, plus an allowance for tracks that may go bad in the future.

17. Next, **badtrk** prompts:

```
Do you want to update this device with the new table? (y/n)
```

Enter **y** and press ⟨Return⟩ to save the changes. To correct any mistakes or otherwise alter the bad track table, enter **n**. Modify the bad track table to contain the desired entries. Enter **q** at the main menu to return to the prompt displayed above, and then enter **y** to upgrade the device with the new table.

18. Now you are prompted:

```
Do you want to attempt to salvage any valid data
on the bad tracks? [may take a long time] (y/n)
```

Generally you should respond **n** to this prompt. However, if you have not made a backup of needed data on the bad tracks and there is no other way of recovering this information, you may want to respond **y**. Keep in mind that this process may take a long time.

19. If you still have valid filesystems, **divvy** prompts for the option to preserve them across the upgrade. As discussed in step 9, this occurs only if you did not increase the maximum number of bad tracks on your hard disk, shrink the size of the UNIX partition using **fdisk**, or change your partitions with **fdisk**:

```
This device contains a valid disk division table.
Your additional (non-root) filesystems
can be preserved across the reinstallation.
Do you want to preserve your additional filesystems? (y/n)
```

Enter **y**, and proceed to step 21.

20. Now **divvy** prompts you for the swap space allocation:

```
There are n 1K blocks in the UNIX area.
Between x and y 1K blocks should be reserved for the swap area.

Enter the swap-space allocation, or press <RETURN>
to get the default allocation of z blocks:
```

Enter your swap space allocation or press ⟨Return⟩ to choose the default value. If you have a large hard disk, you see the prompt:

```
Do you want a separate /u filesystem? (y/n)
```

We recommend that you have a separate filesystem with disks larger than 240 Mbytes, especially if you anticipate having many active users. You should leave at least 30 Mbytes for the root filesystem.

Check the sections on planning your disk layout in the *Installation Guide* for more information on a */u* filesystem.

If you have enough storage for a separate user filesystem, and want to create one, enter **y** and press ⟨Return⟩.

Next, you see:

```
Enter block allocation for the /u file system.
(min to max)
```

In the display, *min* and *max* are replaced with the minimum and maximum number of blocks that can be allocated for the filesystem.

21. Next, you are prompted for precise control over the layout of your filesystem:

```
The layout of the filesystems and swap area is now prepared.

Do you wish to make any manual adjustments to the sizes or
names of the filesystems or swap area before they are created
on the hard disk? (y/n)
```

If you are preserving or recreating any filesystems, you must respond **y**. If you are not recreating or preserving filesystems, respond **n** and proceed to Step 25.

22. A **divvy** division table similar to the one below is displayed, followed by the **divvy** menu:

```
+-------------+-------------+--------+---+-------------+-------------+
| Name        | Type        | New FS | # | First Block | Last Block  |
+-------------+-------------+--------+---+-------------+-------------+
| root        | AFS         |  yes   | 0 |           0 |       47402 |
| swap        | NON FS      |  no    | 1 |       47403 |       50368 |
| u           | AFS         |  no    | 2 |       50369 |       70368 |
|             | NOT USED    |  no    | 3 |           - |           - |
|             | NOT USED    |  no    | 4 |           - |           - |
|             | NOT USED    |  no    | 5 |           - |           - |
| recover     | NON FS      |  no    | 6 |       70369 |       70378 |
| hd0a        | WHOLE DISK  |  no    | 7 |           0 |       70676 |
+-------------+-------------+--------+---+-------------+-------------+

70379 1K blocks for divisions, 298 1K blocks reserved for the system

n[ame]      Name or rename a division.
c[reate]    Create a new file system on this division.
t[ype]      Select or change filesystem type on new filesystems.
p[revent]   Prevent a new file system from being created on ...
s[tart]     Start a division on a different block.
e[nd]       End a division on a different block.
r[estore]   Restore the original division table.

Please enter your choice or 'q' to quit:
```

23. If you are preserving your filesystems, refer to the data that you recorded on earlier in Table 1. Note that the name column may be conspicuously different from your original table, unless you have only one non-root filesystem called */u*. Skipping root and swap, you should rename all of your old filesystems (including */u*), starting with division 2. Enter **n** to begin. You are prompted for the division number (column four):

```
which division? (0 through 6) --
```

Enter **2** and respond with the proper name for that division. Repeat for each of your additional filesystems. If you make an error, you can enter **r** (restore) and all your changes are erased. Do not enter any option other than **n** or **r**. When you are finished renaming divisions, proceed to step 25.

24. If you are recreating your filesystems, use most of the data that you recorded earlier in Table 2 to re-create filesystems at least as large as your original ones, so that your backups fit within the space allocated. To choose a command, enter the first letter of the command, then press ⟨Return⟩.

    Enter **s** to start a division on a different block number. Enter **e** to end a division on a different block number. You can use these two commands to change the size of a partition. Obviously, if you need to increase the size of one division, you must reduce another. If you are increasing the size of your swap space, for example, you could first reduce the *root* filesystem with the e option, then enlarge the swap space with the s option. You could also reduce another filesystem. Note that if any of the divisions overlap, an error message appears when you try to exit and you are returned to the **divvy** menu to correct the situation.

Note

When you recreate a UNIX filesystem, the filesystem type is reset to the default AFS filesystem type. You must enter **t** to change the filesystem type. You are then prompted for the division number. Make certain you use the correct number.

You can enter **r** (restore) option to restore the original partition table if you make a serious mistake and want to start over.

---

25. Enter **q** to quit out of **divvy**.

26. After you quit from **divvy**, you may be prompted for operating system serialization. Follow the prompts.

27. When the system shuts down, you are instructed to insert one of the distribution volumes before starting on your reinitialized hard disk. When you reboot the system, the usual copyright and configuration information is displayed, followed by the A-Z self check.

28. Depending on your media type, you may then be prompted for operating serialization (if you were not already asked to do so), and **fsck** checks your filesystems.

29. You are next prompted to insert additional floppies to complete installation of the Run-Time System.

3.1 Reinstall Packages and Applications

Now that the Run Time System is installed, you should reinstall your packages, applications and device drivers just as you did when you installed them originally.

3.2 Relink the Kernel

If you have special device drivers, you need to link them into the UNIX kernel when you are finished updating. In this case, make sure you install the link kit software when you upgrade your UNIX System. See **custom**(ADM) for information on extracting packages from the UNIX distribution.

When you are finished updating, make sure you are in single user (system maintenance) mode.

3.3 Reestablish Your Second Hard Disk

This section explains how to prepare your second hard disk to function with your upgraded system:

1. Enter single-user mode.

2. If you do not know what **fdisk** partitions your UNIX filesystems are in, run **fdisk**:

   fdisk -f /dev/rhd01

   Enter **1** at the main **fdisk** menu to display the partition table. Write down the numbers of the UNIX partitions, then exit from **fdisk**.

3. Invoke **divvy** with the following command line:

   divvy -b 1 -c 1 -p 1 -v *partition-number*

   In the above command line, substitute the actual partition numbers that you copied from the **fdisk** table. Follow this **divvy** procedure for all of the possible **fdisk** partitions.

   Examine your **divvy** table. You should see a list of **divvy** divisions, which may contain valid filesystems, that do not have names. Only the **divvy** division number and sizes are displayed.

   Next, use the information you recorded in Table 2, "Secondary Hard Disk Filesystems" to rename each of the divisions to their original names. Enter **n** to name the division. If you make a mistake, enter **r** to undo any changes you made. When you have named all of your **divvy** divisions, enter **q** to quit from **divvy**.

4. At this point you are given the choice to retain your changes or exit:

```
i[nstall] Install the division set-up shown
r[eturn] (Return)to the previous menu
e[xit]    Exit without installing a division table

Please enter your choice:
```

If you wish to keep your changes, enter **i**. If you have made a mistake, enter **r** to return to the main **divvy** menu.

5. Finally, use the **mkdev** command to prepare your preserved filesystems for use as described in the next section.

3.4 Prepare Preserved Filesystems for Use

To use the filesystems you preserved, you must create the directories that they are mounted upon. Follow these steps:

1. Use the **mkdev** command to ensure that the filesystems you preserved are correctly set up:

   mkdev fs

   Δ **sysadmsh** users select: Filesystems→Add

2. You see the following:

```
Filesystem Initialization Program
This program performs maintenance tasks required to add
or delete an existing filesystem. Would you like to:
    1. Add a new filesystem to system.
    2. Remove a filesystem.
Select an option or enter q to quit:
```

Select **1**.

3. You are next prompted for the device name:

```
Enter a device name and press (Return) or q to quit:
```

Enter the full pathname of the device from */dev*. For example, to add a filesystem called *u*, you enter */dev/u*.

4. You are now prompted to provide the name of the mount point to be used:

```
Enter a directory name and press (Return) or q to quit:
```

This directory is where the filesystem is mounted. For example, a filesystem called *u* is mounted on the directory */u*.

5. The following is displayed:

```
Reserving slots in lost+found directory ...
When entering multiuser mode:

   1. Always mount filesystem
   2. Never mount filesystem
   3. Prompt before mounting filesystem

Select an option:
```

If you want the filesystem mounted automatically at system startup, enter **1**. If you wish it mounted only by the request of the system administrator, select **2**. If you select **3**, the system prompts you at system startup whether or not you want the filesystem mounted.

6. Next, you see:

```
Do you want to allow users to mount this file system? (y/n)
```

You must respond **y** so that the system backup program can mount and unmount the filesystem as necessary.

7. The following messages are displayed when the process is complete.

```
Updating system files ...
Filesystem has been successfully added.
```

8. Next, you should mount the filesystem using the following command (*/u* is used in this example):

   mount /dev/u /u

   Δ **sysadmsh** users select: Filesystems→Mount

   The filesystem is now ready for use.

3.5 Restore Your Backups

You should restore your backups while still in single user mode if you did not preserve them during the upgrade. Use the Backups→Restore→Full selection of **sysadmsh**(ADM), to restore your backups of non-root filesystems. Do not forget to mount your non-root filesystems before restoring their respective backups.

3.6 Restoring Configuration Files

Having upgraded your UNIX System to Release 3.2 Version 2.0, you may wish to configure the upgraded system to look like the old system. To do this you must restore certain files from your backup media.

3.6.1 /etc/default Files

If you customized any of the files in the */etc/default* directory (for example, */etc/default/filesys*), you should restore them now so that consistency is maintained.

3.6.2 cron Configuration

If you wish to allow the same users to use the **at**(C), **batch**(C), and **crontab**(C) commands, you must remove */usr/lib/cron/cron.** and */usr/lib/cron/at.** then restore the following files (if they exist on your backup):

/usr/lib/cron/cron.allow
/usr/lib/cron/cron.deny
/usr/lib/cron/at.allow
/usr/lib/cron/at.deny

If you want to keep the jobs that were submitted with the **at**(C) and **batch**(C) commands, then you must restore the contents of the

/usr/spool/cron/atjobs directory. If you want to keep the jobs that were submitted with the **crontab**(C) command, then you must restore the contents of the */usr/spool/cron/crontabs* directory.

If you have made any changes to */usr/lib/cron/queuedefs* then restore this as well.

Note

Do not restore the system crontab files (*adm*, *root*, *sys*, and *uucp*).

If you have made any modifications to the batch prototype files, */usr/lib/cron/.proto* and */usr/lib/cron/.proto.b*, then you may wish to restore these too.

3.6.3 UUCP Configuration Files

If you configured your UUCP system previously, you should restore the UUCP configuration files, including:

/usr/lib/uucp/Devconfig†
/usr/lib/uucp/Devices
/usr/lib/uucp/Dialcodes†
/usr/lib/uucp/Dialers†
/usr/lib/uucp/Permissions
/usr/lib/uucp/Poll†
/usr/lib/uucp/Sysfiles†
/usr/lib/uucp/Systems

You need not restore the files marked with a dagger if you did not customize them.

3.6.4 Micnet Configuration Files

If you configured a Micnet network previously, you should restore the Micnet configuration files, including:

/etc/default/micnet
/usr/lib/mail/top

3.6.5 MMDF Configuration Files

If you configured your system to route mail with MMDF, you should restore the following files (the asterisk refers to all the files and subdirectories present in the directory):

/usr/mmdf/mmdftailor
*/usr/mmdf/table/**

3.6.6 User Accounts

If you wish to keep user accounts, then you must restore the following files (the asterisk refers to all the files and subdirectories present in the directory):

/etc/auth/system/default
*/etc/auth/subsystems/**
/etc/group
/etc/passwd
*/tcb/files/auth/?/**

If the users home directories are to be located in a different place, you must change their home directory fields in the */etc/passwd* file. This can be done with the selection Accounts→User→Examine:Identity of the **sysadmsh**(ADM).

Now restore users' home directories; remember that if you restore the directories to a new location, then the users must have search permissions on all components of the path leading to the home directory.

Also, if any users have login shells that are not part of the standard operating system, then restore these as well.

If you have made modifications to the */etc/default/authsh* file, you may wish to restore this as well, so that the same restrictions specified for creating new users in your old system can also be used in the upgraded system.

3.6.7 Terminal Control Database

If you have restored user accounts and wish to keep the contents of the terminal control database, then you must restore */etc/auth/system/ttys* and the contents of the device assignment database */etc/auth/system/devassign*.

3.6.8 Auditing Parameters

If you wish to keep your auditing parameters (such as events audited, and where compaction files are located), then you must restore the */tcb/files/audit/audit_parms* and */tcb/files/audit/parms.saved* files. If you also wish to keep your audit reduction templates, restore the contents of the */tcb/files/auditrparms* directory.

3.6.9 File Control Database

If you made any changes to the file control database (*/etc/auth/system/files*), then you need to make these changes again on your new system. Enhancements have been made to the file control database and the utilities that use it. These utilities do not work reliably with the old file control database, so it is important that you do not simply restore your old file control database instead of making changes to the new version.

Note

If you have only added new entries for extra shells then you do not need to do anything to the new file control database because the enhanced wildcard matching covers any new shell files automatically.

3.7 Configuring the Systems Default Database

If you did not install the C2 Support Level Supplement (unx-167) on your original system, but you did transfer user accounts, then you need to use **sysadmsh**(ADM) to add extra fields to the systems default database */etc/auth/system/default*.

Run **sysadmsh**(ADM) and select Accounts→Defaults→Password. A warning appears at the bottom of the screen stating that some values from the systems default database are missing. The missing value is the default for the new "Password required to login" parameter. **sysadmsh**(ADM) assumes a value of yes in the absence of a real default. Executing this option causes a real default to be written to the systems default database.

Now select Accounts→Defaults→Logins. Again a message appears at the bottom of the screen stating values are missing from the systems default database. The missing value in this case is the new "Time (in seconds) to complete successful login" parameter. In the absence of a real default, **sysadmsh**(ADM) assumes a value of 40 seconds, however, you can change this to a more suitable value. Again, executing this option causes a real default to be written to the database.

The defaults you have chosen for the two new parameters now apply to all users on the system.

You may wish to choose specific values for particular users. These specific values can be set using **sysadmsh**(ADM) with either Accounts→User→Examine:Password for "Password required to login" or Accounts→User→Examine:Logins for "Time (in seconds) to complete successful login."

The upgrade is now complete, and the system can be taken to multiuser mode.

SCO® UNIX® System V/386 Operating System

Release Notes
Release 3.2 Version 2.0

START HERE

1. Assemble your documentation by inserting the tab pages in the places marked with the black-edged pages. Place these *Release Notes* in the binder marked *System Administrator's Reference*.

2. Read "Before Installing Your Software" in these *Release Notes* for information you must know prior to beginning the installation process.

3. Look over the "Compatible Hardware" section in these *Release Notes* to see if your hardware requires special consideration during installation.

4. Use Chapter 2, "Installation Procedure," of the *Installation Guide* to install your UNIX system.

The Santa Cruz Operation, Inc.

Document Version: 3.2v2.0A
Date: 24 September 1990

SCO UNIX System V/386 Operating System Release Notes Release 3.2 Version 2.0

Release Notes
SCO UNIX System V/386
Operating System
Release 3.2 Version 2.0
24 September, 1990

1. Preface

This document contains important information about the SCO UNIX System V/386 Operating System Release 3.2 Version 2.0. These notes are divided into three parts: software notes, Appendix A, "Compatible Hardware" and Appendix B "Using the Korn Shell (ksh)."

The software notes are organized into the following sections:

- Before Installing Your Software
- Special 386 Microprocessor Notes
- Administering Your System
- System Security
- Using Your System
- Using the System Console
- Using Printers
- Using Floppies and Tapes
- Using a Mouse
- Using MS-DOS and OS/2
- Using Networks
- System Configuration and Link Kit Notes
- Internationalization
- Compatibility and Conformance Notes
- Programmers Notes

Note

Please read through the "Before Installing Your Software" section of these *Release Notes* before installing the SCO UNIX System V/386 Operating System Release 3.2 Version 2.0. In addition, pay particular attention to the sections of these *Release Notes* pertaining to peripheral devices such as printers, tape drives, or other hardware you are installing.

Also, please note that certain hardware configuration information included in Appendix A of these notes may be required for successful installation. You should refer to those sections of Appendix A that apply to the hardware in your configuration if you have trouble installing the operating system.

We are always pleased to hear of users' experiences with our products, and welcome recommendations on how they can be improved. We give serious consideration to all written suggestions.

1.1 Conventions Used in These Notes

Utilities and commands are printed in **boldface** type, with the *Reference Guide* section following in parentheses (). Filenames are *italicized.*

1.1.1 SCO UNIX System V/386 Documentation

The SCO UNIX System V/386 Operating System documentation consists of several guides and references. Each has an extended name, such as the following:

SCO UNIX System V/386 Operating System User's Reference

For ease of use, references to the book names within the documentation set are truncated. For example, the above manual is simply referred to as the *User's Reference.*

1.2 Contents of the Distribution

The SCO UNIX System V/386 Operating System Release 3.2 Version 2.0 is distributed on 96tpi and 135tpi diskettes for standard architecture machines.

The distribution on 96tpi diskettes is as follows:

Volume M1
Volumes N1-N5
Volumes B1-B4
Volumes X1-X6

The distribution on 135tpi diskettes is as follows:

Volume M1
Volumes N1-N4
Volumes B1-B4
Volumes X1-X5

The SCO UNIX System V/386 Operating System Release 3.2 Version 2.0 is distributed on 135tpi diskettes for microchannel architecture machines.

Volume M1
Volumes N1-N4
Volumes B1-B4
Volumes X1-X5

The software is grouped into packages as listed in the section, "Packages in This Set".

1.3 Tape Distributions and the Tape Bootstring

The SCO UNIX System V/386 tape distribution configures your tape drive at install time by creating a special bootstring to be used at the prompt:

Boot:

This bootstring is automatically invoked each time the system is started. You do not need to invoke it manually unless you are booting from the floppy drive and you need to restore from backup tapes.

Note

The bootstring option only applies to QIC-02 cartridge tape drives; it does not work for SCSI, QIC-40, or Irwin drives.

This option is intended for use during tape installation and does not replace the functions of **mkdev tape**. If you later run **mkdev tape** to add a cartridge tape drive, you are prompted whether you wish to modify the current tape bootstring, retain it, or remove it entirely.

The bootstring has the general format:

ct=*controller*(*base*,*irq*,*dma*)

Here are some examples:

ct=wangtek(0x338,5,1)
ct=archive(0x330,4,1)

Numbers prefixed with 0x are assumed to be hexadecimal, otherwise numbers are assumed to be decimal. You must also specify the kernel boot device. Thus, a complete boot line might look like the following (the necessary input is shown in bold):

```
Boot: hd(40)unix ct=wangtek(0x338,7,1)
```

If you invoke the tape bootstring manually, note that you must specify **hd(40)unix** or **fd(52)unix**, not just **unix**. The tape bootstring is not checked until the driver is initialized. If the configuration information supplied in the bootstring appears to be invalid (for example, the controller named in the bootstring is not supported) then a warning message is printed, and the tape driver ignores the bootstring and continues to use the default configuration.

1.4 Packages in This Set

The SCO UNIX System V/386 Operating System software is grouped into packages. This makes customizing your system easier because you can use the **custom**(ADM) utility to add or delete groups of programs that have related functions.

For example, if you do not want to use your system for communications, you would not install the UUCP package. If you install a package, then change your mind later, use **custom** to remove that package. **custom** can locate all the files that belong in a certain package and delete them.

Note

At this time, the UUCP package is installed whether or not it is selected. If you are planning to follow the requirements for a C2 system, you must remove the UUCP package using **custom**.

The packages in the SCO UNIX System V/386 Operating System are listed below. The sizes are in 512 byte blocks. Please note that this is a sample listing only. The actual sizes of the packages in your set may differ from those shown here. Use **custom**(ADM) to determine the actual sizes in your set.

UNIX System V/386 Operating System Packages

Package	Size (512 byte blocks)	Use/Contents
ALL	33842	Entire Operating System set
BACKUP	256	System backup and recovery tools
BASE	1406	Basic extended utility set
CSH	116	The C-shell
KSH	282	The K-shell
DOS	368	DOS utilities
EX	408	The ex and vi editors
FILE	480	File manipulation tools
LAYERS	194	System V Layers
LPR	3202	Multiple line printer spooler
MAIL	4944	Electronic mail and micnet
MAN	3764	Operating System Manual Pages
MAPCHAN	144	International character set mapping
MOUSE	68	Mouse and graphic input devices files
SYSADM	868	Additional system administration tools
TPLOT	350	Tplot, Graph, and Spline
UUCP*	2280	Uucp and cu communications utilities
LINK	6816	System V Link Kit files

* The SCO UNIX Operating System Release 3.2 is designed to meet the requirements of the C2 level of "trust" as defined by the *Trusted Computer System Evaluation Criteria*, also known as the "Orange Book." If you plan to follow these guidelines, those software packages marked by an asterisk must not be installed on your system. By not installing these packages, you can ensure that your system operates at a greater level of security. (See "Using Networks" later in these *Release Notes* for more information.)

1.4.1 An Explanation of custom(ADM) Packages

This section explains briefly what each package contains.

BACKUP	The utilities necessary to make copies of the files on your system. With this package, you can back up everything or just the files that have changed recently. You can also restore any part of these backups.
BASE	The basic set of system utilities that are normally installed beyond the runtime system.
CSH	The C-shell command interpreter. An alternative to the Bourne shell, **csh** includes extensive command history features.
KSH	The Korn shell, an alternative command interpreter that supports job control, command line editing, and many other advanced features.
DOS	A set of utilities that allow listing, moving and copying MS-DOS files to or from an MS-DOS diskette or MS-DOS partition.
EX	The **vi**(C) screen-oriented text editor that includes the **ex**(C) editor.
FILE	A number of useful programs for comparing and manipulating files.
LAYERS	A window facility for AT&T Blit terminals. The driver must be linked into the kernel to use this package.
LPR	The printer spooler, which allows multiple print jobs to be queued for printing automatically.

MAIL	This includes a number of subsystems that allow the transport of mail within the system and to other systems. The Micnet local network that permits the transfer of files and execution of remote commands is in this package, as well as the MMDF mail routing system.
MAN	This package includes on-line copies of the System Administration (ADM), Commands (C), File Formats (F), Hardware Dependent(HW), and Miscellaneous (M) manual pages.
MAPCHAN	Allows the use of hardware that uses different character sets. For example, a terminal using one character set can send a file to a printer using another character set without jumbling the results.
MOUSE	The utilities necessary to use mice, graphics bit pads, and so on with the system. Note that the kernel must be rebuilt using **mkdev mouse** before these utilities work correctly.
SYSADM	The **sysadmsh**(ADM) system administrator menu interface and other tools used to administer the system.
TPLOT	Utilities for generating graphs. These programs work on a limited set of hardware.
UUCP	Utilities to connect your system locally or to a worldwide network of other systems.
LINK	The tools necessary to rebuild the system kernel, which is necessary if you are adding capabilities to your system such as tape drives, multiport serial cards, mice, MS-DOS filesystem support, and third-party device drivers.

1.5 Software Support

Software support (in the form of assistance, bug fixes, and so forth) is available to customers who purchased the SCO UNIX System V/386 Operating System for use in the United States and Canada. If you purchased it for use outside of the US or Canada, please contact your distributor or retailer for support information.

Software support is described on an insert in the back of the UNIX documentation.

2. Before Installing Your Software

If a cover letter was included with your distribution, please read it first. The installation instructions are in the *Installation Guide* in the binder marked *SCO UNIX System V/386 Operating System System Administrator's Reference*. Chapter 2, "Installation Procedure," describes the installation process.

2.1 Memory Requirements

Please use the following table to determine the amount of memory you need to run the SCO UNIX System V/386 Operating System:

System	Requirements
Operating System	**minimum**: 2 Mbytes (see Note below) **recommended**: 3 Mbytes **multiuser**: 3 Mbytes or more for optimum performance

---

Note

On some 386 machines, the UNIX kernel cannot recognize the memory between 640 and 1024K because the hardware manufacturer has mapped this 384K of memory to another location for their firmware to use. Because this location varies, the UNIX kernel does not know where to find it.

The operating system requires at least 40 Mbytes of disk space (after formatting) to operate efficiently, including space for user files and swap space. If you plan to have security auditing enabled, much more disk space is necessary to maintain the audit record files.

2.2 Cartridge Tape Configuration

If you have the cartridge tape distribution of SCO UNIX System V/386, you must make certain that your tape drive, whether standard or SCSI, is configured so that SCO UNIX System V/386 can recognize and access it automatically. The configuration requirements for standard architecture machines are as follows:

Cartridge Tape Drive Name	Address	DMA	Interrupt
Standard	0x0338-0x033C	1	5
Mountain	0x0200-0x0203	1	5
IBM 6157	0x3120-0x312F	2	6
Everex/Archive	0x0300-0x0301	1	3
Tecmar	0x0300-0x0301	3	5

SCSI Cartridge Tape	
Host Adapter	0
ID	2
LUN	0

If you do not configure your drive as described above, you are required to reconfigure the system during installation and restart before continuing the installation. For information on supported hardware, see "Tape Drive/Controller Combinations" in Appendix A of these notes. For more information on configuring tape drives, see the "Using Floppy Disks and Tape Drives" chapter of the *System Administrator's Guide*.

2.3 Hard Disk Interleave

Starting with this release of SCO UNIX System V/386, 1:1 is now the preferred interleave for disk controllers that support it. Most modern disk controllers support a 1:1 interleave, but older, single-buffered MFM controllers do not. A controller's fastest/best interleave is one of its most widely published specifications; check your documentation. (Typically, a 3:1 is suggested for MS-DOS.) However, if you are unsure and unwilling to experiment, we suggest that you format your drive using a 3:1 interleave. Although with an appropriate controller a 1:1 disk interleave can substantially improve I/O performance, an inappropriate disk/controller combination formatted 1:1 can reduce performance by a factor of 4 or worse. (Operations that would take 15 seconds on a drive with proper interleave could take as long as a minute.) To determine if you have an interleave problem, first perform a minimal installation, then enter the command:

dd if=/dev/root of=/dev/null bs=1k count=4096

If this command takes longer than 45 seconds to complete, you could have a disk interleave problem and you may wish to reformat and reinstall before proceeding further. (However, it is also possible that your particular disk/controller combination is intrinsically slow, even when formatted at the proper interleave.)

Most new disk/controller pairs arrive already formatted at the proper interleave (3:1 for MS-DOS). If reformatting is necessary, it must be performed using the setup disk that came with your system, or a DOS formatting program.

Note that some disk/controller pairs, notably SCSI, have their own cylinder/head/sector translation logic, and will probably ignore any reformat command sent to them.

2.3.1 Future Controllers

At this time the slowest 386-based CPU, the 80386SX, can barely keep up with the fastest 1:1 double-buffered controllers. As transfer speeds surpass 10 Mbytes/sec, slower CPUs may or may not be able to keep up with the data transfer rate. Whether a problem occurs also depends upon the presence of on-chip cache and the amount of buffering on the controller. We are not aware of any specific CPU/controller combinations that cause a problem. If you find yourself with a fast ESDI controller, such as one with a 15 Mbytes/sec transfer rate, a machine based on a comparatively slow CPU such as an 80386SX, and find that disk I/O performance is far worse than you would have expected with a slow disk, your only software solution is to revert to a 2:1 or 3:1 disk interleave.

2.4 If the Installation Stops at H6

When you first start the installation from the N volumes, information is displayed concerning your hardware configuration, followed by a series of diagnostic letters. These letters are displayed and overwritten quickly, so they are not normally visible unless the startup process halts unexpectedly. If the screen display stops at H6 and you have an EGA adapter installed, it is likely that the adapter card is improperly installed. Shut the computer off and check your EGA card documentation. If your graphics card has the AUTOSWITCH feature, it must be disabled to install SCO UNIX System V/386. The documentation should indicate how to disable AUTOSWITCH, which is typically controlled by setting a small switch on the card to the off position. You can then restart the installation from the beginning. If the problem persists, contact the provider of your software for support.

Another symptom of this problem is observed when the boot prompt is displayed:

```
Boot
:
```

If you are using an AUTOSWITCH EGA card and it normally displays in high resolution (EGA mode), and this prompt is displayed in low resolution (CGA mode), you should reset the machine until the high-resolution display appears. (Standard CGA text is low-resolution; dark lines are visible through the characters. EGA text is high resolution; no dark lines are visible.)

2.5 Additional Installation Notes

Read the *Release Notes* and *Installation Guide* and make sure you understand the installation process completely before installing the product. In addition, take note of the following points:

- If you need to stop the installation process for any reason, press the computer "reset" button or turn the power off and on and start the process again from the beginning rather than trying to proceed from the stopping point. Do not abort the installation process by using the ⟨DEL⟩ or ⟨CTL⟩ keys, unless a message appears that explicitly tells you to abort in this way.

- The UNIX system and MS-DOS must boot from the first drive, known as physical drive 0 (hard disk or diskette). Keep this in mind when planning for extra hardware.

  If you are using an ESDI or SCSI disk for your root hard disk, see "ESDI Notes" and "SCSI Notes" later in these *Release Notes*.

- If you are planning to set up your hard disk to accommodate other operating systems, such as MS-DOS, see the "Using MS-DOS and OS/2" chapter of the *System Administrator's Guide*. **Note that only MS-DOS 3.3 or earlier releases are supported for use with SCO UNIX System V/386.**

2.6 lost+found Entries

As shipped, the directory *lost+found* has zero entries. This means that files cannot be recovered by **fsck**(ADM) unless these entries are created. Immediately after installing SCO UNIX System V/386, log in as *root* and enter the following commands:

```
cd lost+found
for I in 0 1 2 3 4 5 6 7 8 9
The system responds with a ">" prompt
do
for J in 0 1 2 3 4 5 6 7 8 9
do
touch ./$I$J
done
done
After a few moments the normal system prompt "#" returns
rm *
```

This creates the necessary entries and leaves the directory empty once again.

2.7 ESDI Notes

This section provides general information for SMS OMTI 8620 and 8627 disk drive controllers. This product supports one SMS OMTI 8620 or 8627 controller with up to two drives attached to it. Either or both drives may have an ESDI or ST506 drive interface. An OMTI-specific device driver supplements the Western Digital-specific device driver.

Note that ESDI disks attached to controllers such as WD1007 use the standard ST506 interface, and thus appear to the system as ST506 disks.

2.8 SCSI Notes

For standard architecture machines, SCO UNIX System V/386 runs on industry standard 386-based systems with an Adaptec AHA-154x SCSI host adapter (or AHA-1640 for microchannel machines), or an IBM SCSI adapter, or Western Digital WD7000 instead of, or in addition to, a standard ST506 (or compatible) disk controller. If you are running on industry standard (AT) architecture, and have only SCSI disks installed on your system, you must run your computer's setup program and set the computer up for operation *without* a hard disk before installing the operating system. This forces the computer to recognize the SCSI adapter.

See "Disk Controllers and Host Adapters" and "SCSI Guidelines" in Appendix A for information on hardware configuration.

2.8.1 Using a Non-Adaptec Adapter

Note that when adding a SCSI device to either an ST506 or EDSI controller (MC architecture machines only), the system assumes by default that you are using an Adaptec SCSI controller.

To use a different adapter, you must use a file editor to make simple modifications to:

- the configuration file */etc/conf/sdevice.d/ad*
- the corresponding file for the adapter you are using
- the file */etc/conf/cf.d/mscsi*
- the file */etc/conf/cf.d/mdevice*.

These instructions assume that you use the **vi**(C) editor, but you can use the editor you are accustomed to. (Note that the EX package must be installed to use the **vi** editor.) To make the modification, follow this procedure:

1. Log in as *root* and enter system maintenance mode.

2. Execute the appropriate **mkdev** script or **sysadmsh** selection. Answer **n** when you are prompted as to whether you wish to re-link the kernel.

3. Enter the following command:

   cd /etc/conf/sdevice.d

4. Edit the file *ad*, as follows:

   vi ad

5. The file contains one line. Change the "Y" in the second field to "N":

   ```
   ad    Y    1    5    1    12    230    232    0    0
   ```

   When you have checked the line to make sure it is correct, save the file.

6. Edit the corresponding file for the adapter you are using, (for example the file *hf* for the IBM SCSI adapter, or the file *wdha* for the WD7000 SCSI adapter) as follows:

   vi hf

7. The file contains one line, similar to that in step 3. Change the "N" in the second field to "Y". When you have checked to make sure that the line is correct, save the file.

8. Enter the following command:

   cd /etc/conf/cf.d

9. Edit the file *mscsi*, as follows:

   vi mscsi

10. The file contains one line for each SCSI device configured into the system. Ensure that the first field in each line (the adapter type) matches the type of adapter that you are using. When you are sure that all the entries are correct, save the file.

11. Edit the file *mdevice*, as follows:

    vi mdevice

12. Search for "ad", which corresponds to the entry for the Adaptec controller. It will look similar to the line below:

```
ad    iI    iHroCc    aha    34    34    1    2    5
```

13. If the third field of your entry reads "iHroCc" or similar (that is, it contains an "r"), remove the "r". Do not save the file at this point.

14. Search for the entry corresponding to the controller which you are using (for instance, "hf" for an IBM SCSI adapter). The entry will look similar to the example shown in step 12.

15. If the third field reads "iHoCc" or similar (that is, it does not contain an "r"), add an "r" to this field. When you have checked the line to make sure the line is correct, save the file.

16. Enter the following command to re-link the kernel:

    ./link_unix

17. Reboot your machine.

18. The device has been added to the system. Note that with some SCSI devices (for instance, hard disks), the **mkdev** command must be invoked again to complete device configuration. Refer to the appropriate section of the *System Administrator's Guide* for details.

2.8.2 Using a Second SCSI Adapter

If you are adding a second SCSI adapter, ensure that your second adapter is of the same type as your first adapter.

When using **sysadmsh**(ADM) to add a disk or tape device on a second SCSI adapter (controller ID = 6, host adapter = 1, LUN = 0), the following message is displayed:

```
warning: mscsi- device ad; not enough sdevice entries
```

To add the devices successfully, you must use a file editor to make minor modifications to two configuration files: */etc/conf/cf.d/mdevice* and */etc/conf/sdevice.d/ad*. These instructions assume that you use the **vi**(C) editor, but you can use the editor you are accustomed to. (Note that the EX package must be installed to use the **vi** editor.) To make the modifications, follow this procedure:

1. Log in as *root* and enter the following command:

   cd /etc/conf/cf.d

2. To find out which SCSI adapter driver is being used, enter the command:

   cat mscsi

   The first field of the last line in this file shows which adapter is in use.

3. Edit the file *mdevice*, as follows:

   vi mdevice

4. Search for the entry which corresponds to the adapter you are using. For example, the entry for "ad" looks similar to the line below:

   ```
   ad    iI    iHroCc      aha    34    34    1    2    5
   ```

   The entry for "hf" looks similar to the line below:

   ```
   hf    iI    iHroCc      hf     36    36    1    2    -1
   ```

5. If the third field of this entry reads "iHroCc", continue to the next step. If it reads "iHoCc", then the adapter corresponding to that entry is not in use.

6. Change the third field of this entry from "iHroCc" to read "iHrCc" by removing the "o".

7. Save the file when the change is correct.

8. Enter the following command:

   cd /etc/conf/sdevice.d

9. Edit the appropriate adapter file, as shown for "ad" in the following command:

   vi ad

10. The file contains one line. You must add another line similar to the one already present. Add the following new line to the file, separating each field with a ⟨TAB⟩:

    ```
    ad    Y    1    6    1    12    230    232    0    0
    ```

    The example shown above is for the "ad" adapter.Please refer to the **sdevice**(F) manual page and your hardware setup information to determine your specific adapter parameters. When you have checked the line to make sure that the line is correct, save the file.

11. Execute the proper **sysadmsh** selection to add the device successfully.

2.9 Saving an Existing XENIX or UNIX Partition

If you have an existing XENIX or UNIX partition on your hard disk that you wish to preserve, you must select the "Fully Configurable Disk Initialization" during the installation or it is wiped out.

2.10 If You Plan to Install TCP/IP

If you are planning to install TCP/IP (release 1.1.2 or earlier) on your system, the TCP/IP installation replaces the */bin/login* binary. This causes the system not to accept further logins. Immediately after exiting **custom**(ADM) from installing TCP/IP, enter the following command:

mv /usr/lib/custom/save/login /bin/login

The system then behaves normally.

2.11 If You Install VP/ix

When installed, VP/ix adds files to the directory */usr/lib/mkuser* with permissions that are not correct. After installing VP/ix, enter the following command:

/tcb/bin/integrity -e > /tmp/int.out

Examine the output in */tmp/int.out* for errors that relate to the */usr/lib/mkuser* directory. Use the **chmod**(C) and/or **chown**(C) commands to set the proper permissions on the files reported by **integrity**(ADM).

2.12 If You Install the Development System

If you install the SCO UNIX System V/386 Development System 3.2.0 on your system, certain files are overwritten that must be replaced from your Operating System distribution. After installing the Development System, do the following:

1. Enter **custom** and press ⟨Return⟩ to select Install.

2. When the product window opens, use the Down Arrow key to highlight the SCO System V Operating System and press ⟨Return⟩.

3. Use the Left Arrow key to highlight Service Components and press ⟨Return⟩.

4. When the service component window opens, press ⟨Return⟩ to select SCO System V Runtime System.

5. Use the Left Arrow key to highlight Files and press ⟨Return⟩.

6. When the package window opens, press ⟨Return⟩ to select Entire RTS with perm lists.

7. When the file list appears, press ⟨F5⟩ key. What you need to do is first search for a file name, then mark it with the ⟨Space⟩. Do this for each of the following files, one at a time:

   /usr/bin/infocmp
   /usr/bin/mcs
   /usr/bin/tic

   You are prompted to insert the necessary distribution volumes.

8. After these files are restored, you are returned to the top of the **custom** menu. Press ⟨Return⟩ to select Install and repeat steps 2-3, selecting SCO System V Extended Utilities this time.

9. Use the Left Arrow key to highlight Files and press ⟨Return⟩.

10. When the package window opens, press ⟨Return⟩ to select Entire Extended Utilities.

11. When the file list appears, follow the same procedure as in step 7 to restore the following files:

    /bin/ar
    /bin/ld
    /lib/coff/ar
    /lib/coff/ld
    /usr/bin/what

    When the process is complete, all the necessary files have been restored.

2.13 If You Remove the Development System

If you remove the Development System at a later time, the file */etc/perms/dsmd* is removed. This file is actually part of the Operating System distribution and must be present on the system. After removing the Development System, use **custom** to restore the */etc/perms/dsmd* file from the SCO System V Runtime System as described in "If You Install the Development System."

2.14 If You Install Portfolio

After installing Portfolio successfully, it runs properly, but after relinking the UNIX kernel, running Portfolio results in the error message:

```
Can't open pseudo tty (errorno = 2)
```

This is because the device */dev/mvwdev* is not preserved when the kernel is relinked. To correct this problem, edit the entries **mvwc** and **mvws** in the file */etc/conf/cf.d/mdevice*. Using a file editor, change the 3rd field in these entries from "ico" to "icor". The "r" indicates the entry is required. As long as the required flag is added to the **mvwc** and **mvws**, subsequent relinking will not remove */dev/mvwdev*.

If you have already relinked the kernel and lost the device, log in as *root* and enter the following commands:

cd /etc/conf/cf.d
./configure -j mvwc

This command will return the major number of the mvw controlling device. Use the number obtained as the major number, *major*, in the following command:

mknod /dev/mvwdev c *major* **255**

3. Special 386 Microprocessor Notes

The following are notes that apply to known problems with 80386 chips.

3.1 Intel 387 Coprocessor Problems

Because of design defects in Intel's 80386 chip (B1 stepping), the Intel 80387 math coprocessor may not operate correctly in some computers. The problem causes the CPU to hang when DMA/paging/coprocessor accesses are occurring. You can work around this problem by changing the tunable kernel parameter DO387CR3. You can do this using the **configure**(ADM) utility. See the "Tuning System Performance" chapter in the *System Administrator's Guide*.

You may replace the 386 chip with a newer release of the 386 chip (a D-step part), or you can bypass the 387 chip by adding the **ignorefpu** keyword in your boot command as follows:

```
Boot
: unix ignorefpu
```

This means that the operating system will not use the 387 chip, but you need not remove it physically; the coprocessor is still usable from MS-DOS. To bypass the 387 chip automatically every time you boot your system, add the **ignorefpu** keyword to the */etc/default/boot* file. See **boot**(HW) for more information.

There are hardware devices available to cope with this problem. See *Ironwood Electronics* in the *SCO System V Directory*.

3.2 Using an Intel Inboard in High-Speed Mode

The Intel Inboard is a plug-in card for a 286 AT standard architecture machine that replaces the 286 CPU with a 386 CPU. To use the high-speed, cache mode of the Intel Inboard, add the keyword **inboard** to the boot line:

```
Boot
: unix inboard
```

To use this mode of the inboard every time you boot the system, add the **inboard** keyword to the */etc/default/boot* file. This parameter is passed to the kernel and sets the mode of the board. See **boot**(HW) for more information.

4. Administering Your System

This section discusses features that concern system administration.

4.1 Adding a New Filesystem

On page 10-21 of the *System Administrator's Guide*, the procedure for adding a new filesystem contains an error. Step 9, which includes **chmod** and **chgrp** commands, should be as follows:

chmod 755 /x;chgrp auth /x

The commands alter the permissions of the mount point, not the device as documented.

4.2 Migrating XENIX Accounts: addxusers(ADM)

addxusers(ADM) is a tool for moving XENIX system accounts to SCO UNIX System V/386. See the manual page **addxusers**(ADM) in your *Replacement Manual Pages* for complete instructions concerning the use of this command.

4.3 Performance Enhancements in This Release

System performance was greatly enhanced in this release. Two hard disk interleaves are now supported: 1:1 and 3:1 (see the "Hard Disk Interleave" section earlier in these *Release Notes*). In addition, two tunable disk efficiency schemes were added: namei caching and scatter-gather I/O. These schemes are discussed fully in the "Tuning System Performance" chapter of the *System Administrator's Guide*.

4.4 System Administrator's Guide

The *System Administrator's Guide* was enhanced and expanded for this release. The chapters have been reorganized for ease of use, three chapters have been completely rewritten, and one has been added:

- "Introduction" includes new information on site planning, tasking, and user education. The information contained in this section will help you to set up your system and perform maintenance on a regular basis.

- "Troubleshooting Your System" now contains all the information related to solving problems. If you have any problems installing or maintaining peripherals, or with error messages displayed on the console, this chapter will help you to solve them.

- "Setting Up Electronic Mail" was rewritten and new material added to allow you to take full advantage of the configurability of the MMDF mail router.

- "Using the Audit Subsystem" explains in detail how to set up and use auditing to enhance the security of your system.

4.5 Shared Library Compatibility

There is a discrepancy between the network services shared library distributed with SCO UNIX System V/386 and other UNIX Systems. Programs compiled using the network services shared library from other UNIX Systems will not work under SCO UNIX System V/386 and vice versa. To solve this problem, two sets of the shared library are supplied. The shell script */usr/bin/fixshlib* can alter an imported binary to call the alternate shared library rather than the default SCO shared library. The script is invoked as follows:

fixshlib *progname*

where *progname* is the name of the imported UNIX binary. The program then uses the alternate network services shared library.

4.6 SCO ISAM Runtime Support

This release contains the SCO ISAM utility **isverify**(M). See the manual page **isverify**(M) in your *Replacement Manual Pages* for full details of this utility. **isverify** is used to validate and repair ISAM files associated with applications created with the SCO ISAM Development System.

4.7 CD-ROM Support

This release supports the CD-ROM filesystem formats "High-Sierra" and ISO9660. (See Appendix A for a list of supported CD-ROM drives.) In order to use the CD-ROM feature, you must add support for the filesystem to the kernel using the **mkdev high-sierra** command.

The **mkdev cdrom** command (undocumented in **mkdev**(ADM)) adds a CD-ROM drive, and should be run in multi-user mode after you have completed the UNIX installation. When you run **mkdev cdrom**, you are asked to provide the SCSI controller-id and logical unit number (LUN) of each of the devices that you are installing.

mkdev cdrom also asks which SCSI host adaptor the devices are connected to. Normally only one SCSI host adaptor is installed, and this is configured as host adaptor 0, therefore you should normally specify host adaptor 0 for all SCSI devices on your system. Finally, you are asked whether you wish to relink a UNIX kernel which includes the device drivers that have just been installed. If you wish to make use of the devices you should answer "yes" to this question.

Note

To bring the CD-ROM drive online, you must insert a disk. If you attempt to bring up the drive without inserting a disk, the message "cannot open" is displayed.

4.7.1 CD-ROM Device Names

The CD-ROM driver supports both block and raw device access to the CD-ROM drive.

The device names for CD-ROM drive 0 are:

/dev/cd0	block device
/dev/rcd0	raw device

The manual page for **cdrom**(HW) in your *Replacement Manual Pages* contains more information about this device driver.

4.7.2 Using High Sierra and ISO9660 filesystems

The ISO9660/High Sierra CD-ROM filesystem can be mounted as a read-only filesystem under SCO UNIX System V/386. The filesystem allows access to files which are described by the primary volume descriptor on the CD-ROM. Access to files described by secondary volume descriptors is not supported. The current release of software provides a level of support which is similar to that provided by the Microsoft CD-ROM Extensions for MS-DOS.

Under SCO UNIX System V/386 all files on the CD-ROM appear to have access permissions of 555 (that is, files are readable and executable by all users and are not writable). Filesystems containing extended attribute records are supported, but the record format information and the file access permissions in the extended attribute record are not used.

CD-ROM filesystems are mounted with the **mount**(ADM) command. All utilities behave as expected. You can traverse the filesystem with the **cd** command, and so forth, including copying files from the CD-ROM filesystem to other filesystems with such utilities as **cp**(C) and **tar**(C).

4.8 Backups Using sysadmsh(ADM)

The backup level scheme was revised for this release. The backup **schedule**(ADM) file now uses levels 0, 1, 2, and 3 instead of 0, 1, 8, and 9. This was done to make the levels easier to understand without breaking backwards compatibility. Note that the **schedule**(ADM) manual page does not reflect this change.

4.9 Setting System-Wide Variables: initscript(ADM)

Environment variables such as **HZ** are now set in the file */etc/initscript* rather than in several different locations. See **initscript**(ADM) for more information.

4.10 crontab -u Option

A **-u** option was added that allows the cron administrator to use the **-r** and **-l** options as if they are the owner of the crontab file. The crontab file cannot be replaced, however, meaning that the option can only be used with **-r** (remove) or **-l** (list) options.

4.11 The assign(C) Command

If you wish to use the **assign**(C) utility (which allows the system administrator to assign a device to a specific person and disallow access to anyone else) some modifications are necessary. Follow these steps:

1. Change ownership of the */dev* directory (user *bin* has to be able to write a lock file in */dev*):

   chown bin /dev

2. Change the ownership of the specific *devices* to *asg* and change their mode to disallow access to others:

   cd /dev
   chown asg *devices*
   chmod 660 *devices*

Make certain you invoke **assign** or **deassign** from a terminal where the output of both **tty**(C) and **who**(C) indicate the same device. If this is not the case, **assign** gives a "No tty" error. This often happens on *tty01* (which is also usually known as */dev/syscon*).

4.12 Using the 8-Bit Character Set

To use 8-bit character sets and have them recognized by UNIX utilities, you must configure the 8859 character set in the file */etc/default/lang*. The existing line must be changed to appear as follows:

```
LANG=english_us.8859
```

The default setting of ASCII configures the standard 7-bit character set. Individual users can also set the **LANG** variable in their *.profile* or *.cshrc* files. See **locale**(M) and **environ**(M) for more information.

4.13 Making Tape Backups with cpio

A **-K** *volumesize* option was added to **cpio**(C) that specifies the size of the media volume. You must provide the size in 1024 Kbyte blocks. Note that you must provide the appropriate volume size. If you specify an incorrect size, the command executes without error, but **cpio** generates the message "out of sync: bad magic" when the volume is read.

The **sysadmsh** Backups→Create→Unscheduled selection is the easiest way to back up a filesystem. The Unscheduled backup is simply a complete copy of the filesystem.

If you wish to use the **cpio** command line, use it in a form similar to this example:

```
find . -depth -print | cpio -ocvB -C10240 -O/dev/rct0 -K60000
```

This example is for DC600A tapes. The volume size for a 1.2 Mbyte floppy disk is 1200. You do not need to specify tape length when restoring.

4.14 Single-User Mode

The *root* group ID is incorrectly set to 0 (root), instead of 1 (other) when in single-user mode. To correct this, execute the **newgrp other** command immediately after entering single-user mode.

4.15 Using the sfmt Command

The manual page for **sfmt**(ADM) is incorrect. **sfmt** should only be executed from the "Boot:" prompt.

4.16 POSIX Features

This section reviews the most significant features that conform to POSIX 1003 and FIPS PUB 151-1. (See the "Compatibility and Conformance Notes" later in these notes for a detailed account of conformance issues.)

4.16.1 Restricted chown

This is implemented via the **chown** system kernel authorization. If a site wishes to be FIPS compliant, the system administrator must remove the **chown** kernel authorization from the set of authorizations for default users as described in "Default Account Configuration" in the "Administering User Accounts" chapter of the *System Administrator's Guide*.

4.16.2 Parent and Group ID on Files

This is governed by setting the GID bit on directories. If the GID bit is set on a directory, files created within that directory will have the group ID of the parent directory instead of that of the user or creating process. Without this bit, the GID of a new file is set to that of the creating process or user. See "Setting Directory GID Bit" in the "Managing Filesystems" chapter of the *System Administrator's Guide*.

4.16.3 Multiple Groups

The system allows any user to be a member of up to eight groups simultaneously. This is defined by the kernel parameter NGROUPS, which can be altered using **configure**(ADM). See "Adding/Changing Groups" in the "Administering User Accounts" chapter of the *System Administrator's Guide*. Note that IPC processes (such as shared memory and semaphores) only pay attention to the primary group in this release.

4.16.4 Archiving Tools

POSIX **pcpio**(C), **ptar**(C), and **pax**(C) are included.

4.16.5 Job Control

Job control is fully implemented in this release and all job control features are accessible though the **ksh**(C) shell. **sh**(C) and **csh**(C) do not support job control at this time.

4.16.6 Filename Truncation

POSIX FIPS requires the system to return an error when asked to create files with names longer than 14 characters (via **open**(), **creat**(), and so on). It is possible to modify this behavior to truncate long filenames silently by using the **configure**(ADM) command to change the ETRUNC tunable parameter.

4.17 MMDF Mail Router

This release includes the Multichannel Memorandum Distribution Facility (MMDFII b, update #32) as a mail router. The operating system uses MMDF to route mail locally and over Micnet, UUCP, or other networks that provide MMDF support.

Note

If you only use local mail, (that is, just on your one machine) then MMDF is automatically configured. When the system is installed, you are prompted to provide a name for your machine. Legal names consist of a maximum of 15 lowercase letters, numbers, hyphens, or underscores. Eight-bit characters are not allowed.

If you wish to configure networking or import alias and routing tables from another release, you must read the "Setting Up Electronic Mail" chapter of the *System Administrator's Guide* for instructions.

4.17.1 Micnet and XENIX Systems

If you are using Micnet between XENIX systems that do not use MMDF and SCO UNIX System V/386 systems, an additional configuration step is required. This is because sending mail to a user on an SCO UNIX System V/386 system from a non-MMDF XENIX machine results in a return address of user@mmdf-host.UUCP, falsely indicating that the sender was on the same system as the receiver.

The solution is the following:

1. Create a global alias file on the XENIX system.
2. Use **netutil** to administer the alias file on the XENIX system, including copying to floppy.
3. Import it to the SCO UNIX System V/386 system and convert it using **mmdfalias**(ADM).

This process must be repeated when the alias file on the non-MMDF XENIX system is updated.

4.17.2 Known Limitations

The following are known problems with MMDF:

- **checkup**(ADM) currently reports incorrect file permissions.
- Although MMDF uses a **dbm**(S)-like set of routines to handle database files in */usr/mmdf/table*, the format is not compatible with the standard **dbm**(S) routines.
- MMBXPREF and MMBXSUFF currently do not work for changing the delimiters in a mailbox.
- If the MMBXNAME parameter is used to specify the name of a mailbox file, you have to create it by hand.
- Global aliases are limited in size. The RHS (Right Hand Side) of a global alias (one defined in *alias.user* and *alias.list* in */usr/mmdf/table*) cannot exceed approximately 8 Kbytes when fully expanded.

- If the "metoo" mode is not set in **mail**(C), the **submit**(ADM) program attempts to make sure that the sender does not get a copy of a piece of mail just because an alias that the user is on is one of the recipients. This code fails if there is an alias that just maps an alias for that one user (for example, *fredf* is mapped to *fred*, the user's real username). This results in the expansion of an alias to a null list that generates an error message.

- Setting "ap=822" on a channel results in the generation of some illegal addresses (for example, local mail to fred and barney may result in "fred barney@sco.COM").

- Incoming UUCP mail may get the local domain address added when inappropriate (for example, mail from "machine!root" may become "machine!root@sco.COM").

- Setting a size limit for MMDF log files does not work properly.

- The first time the system goes to init level 2, the system name used is *scosysv* instead of whatever was defined at install time. When the system is rebooted, the new name is used thereafter.

4.18 SCO Professional 2.0.0 Fix (386 version)

There is a problem with running SCO Professional 2.0.0 (386 version) under SCO UNIX System V/386. This is a known problem in SCO Professional that was fixed in release 2.0.1. To run SCO Professional it is necessary to apply the patch described here.

Once you have installed SCO Professional, log in as *root* and enter the following commands:

```
cd /usr/lib/pro
cp calc calc.00
/etc/_fst -w calc
0x17:3c014?w 0x16eb
$q
sync
```

You can then invoke Professional as usual.

4.19 Floating-Point Emulation

Two floating point emulators are provided in the product. The default emulator, called */etc/emulator.ms*, is linked to */etc/emulator*. The AT&T emulator, */etc/emulator.att*, is not functional and should not be substituted for the Microsoft emulator.

4.20 uname -X

uname(C) has a new option, **-X**, that displays all the information that XENIX **uname**(C) displays, along with some additional information. The following is a sample output:

```
System = unix
Node = scosysv
Release = 3.2v2.0
KernelID = 90/03/06
Machine = i80386
BusType = AT
Serial = scoxxxxxx
Users = unlim
OEM# = 0
Origin# = 1
NumCPU = 1
```

4.21 boot(HW) systty Option

The option **systty=scrn** documented on the **boot**(HW) manual page is incorrect. The correct string is **systty=cn**.

4.22 fsck(ADM) Options

The **fsck**(ADM) manual page lists the **-rr** option, but does not specify that this option is only for XENIX filesystems. An undocumented option, **-b** is similar and works on UNIX filesystems.

4.23 enable(C) Documentation Error

The **enable**(C) manual page lists an incorrect pathname for the file */etc/conf/cf.d/init.base*. Additional information on configuration files can be found in the **inittab**(F) manual page.

4.24 Changing root Shell to Korn Shell

Although it is possible to change the login shell for *root*, it is not recommended. Using any shell other than the standard Bourne shell as the *root* login shell can cause the system to malfunction.

To change the *root* login shell to **ksh**, do the following:

1. Bring the system down to single-user mode.

2. Edit */etc/passwd* and add **/bin/ksh** at the end the *root* entry.

3. Edit */.profile* and make the following changes:

   Change SHELL=/bin/sh to SHELL=/bin/ksh

   Add the following lines at the end of the file:

   ```
   ENV=/.kshrc; export ENV
   set -m
   stty susp '^Z'
   ```

   Close the file when the changes are complete.

4. Enter the following command:

   cp /usr/lib/mkuser/ksh/kshrc /.kshrc

4.25 gettydefs(F): AUTO Keyword

The AUTO keyword documented in the **gettydefs**(F) manual page and the "Administering Terminals" chapter of the *System Administrator's Guide* is not supported. Use of this method for executing an application instead of the **login**(M) program is not necessary. To run an application or other program automatically at a terminal, you should place an appropriate line in the file */etc/inittab*. Follow the guidelines in "Changing the gettydefs File" in the "Administering Terminals" chapter of the *System Administrator's Guide*. The line should contain the full pathname of the program to be run. To make this modification permanent (that is, preserved when the kernel is rebuilt) you should make the same modifications to the */etc/conf/cf.d/init.base* file. In addition, the tty must be disabled (see **disable**(C)) for the program to work.

4.26 swconfig(C)

swconfig(C) sometimes reports products as being "installed" and "partly installed" simultaneously.

4.27 xdumpdir(ADM) and Alphabetized List

The Alphabetized List in the *User's Reference* mistakenly lists the **xdumpdir**(ADM) command as **dumpdir**(C).

4.28 useshell Error Messages

If you see these messages while adding or modifying a user account:

```
useshell: file access control database inconsistency
useshell: not authorized
```

This means there is a problem with the permissions or ownership of a database file located in */usr/lib/mkuser*. You should run the **integrity**(ADM) utility as *root*:

/tcb/bin/integrity -m -e > int.report

Print the file *int.report* and examine it. **integrity** reports files and directories that are missing or have incorrect permissions or ownership. There may be many messages that have no bearing on your problem; pay particular attention to messages about files in */usr/lib/mkuser*. Here is a sample message generated by **integrity**:

```
/usr/lib/mkuser/csh (wildcard entry 216) is wrong.
    Owner is bin, should be root.
    Mode is 0700, should be 0750.
```

In this case, these errors would be corrected with the following commands:

chmod 750 /usr/lib/mkuser/csh
chown root /usr/lib/mkuser/csh

For more details, see "System File Integrity Checking: integrity(ADM)" in the "Maintaining System Security" chapter of the *System Administrator's Guide*.

5. System Security

SCO UNIX System V/386 is designed to meet the requirements of the C2 level of "trust" as defined by the Department of Defense's *Trusted Computer System Evaluation Criteria*, also known as the "Orange Book." To implement the principles of trusted operation, a number of modifications were made that greatly exceed the security features of most UNIX Systems. Each aspect of operation is simplified by selection, available through the **sysadmsh**(ADM) menu interface.

If you do not plan to follow the C2 guidelines, it is possible to configure more traditional UNIX system behavior at installation time or later using the **sysadmsh**.

All aspects of security configuration and adding users is covered in the "Administering User Accounts" chapter of the *System Administrator's Guide*. The C2 requirements for security documentation are satisfied by the "Maintaining System Security" and "Using the Audit Subsystem" chapters of the *System Administrator's Guide* and "Using a Trusted System" of the *User's Guide*.

5.1 Security Enhancements in This Release

Several enhancements relating to security were added to this release:

- automatic security database checking and recovery mechanism for critical database files (**tcbck**(ADM))
- accounts can now be without passwords (as opposed to having ⟨Return⟩ as a password)
- subsystem authorizations are fully functional
- **goodpw**(ADM) is now included
- users can generate audit reports on their own activities when granted the **audittrail** subsystem authorization
- dialup passwords now function properly
- the system administrator can now create user accounts that share a single home directory

These features are fully described in "Administering User Accounts," "Maintaining System Security," and "Using the Audit Subsystem" in the *System Administrator's Guide*. **addxusers**(ADM) explains how to migrate XENIX */etc/passwd* files to your system.

5.2 Editing the /etc/passwd File

You should not edit */etc/passwd* with a text editor. On other UNIX Systems, this is a common but non-trusted way of adding and removing users from a system. On SCO UNIX System V/386 this can generate error messages and cause the system not to accept any further logins. Changing a user's existing entry in */etc/passwd* is equally unreliable. Only some of the fields in */etc/passwd* (and */etc/group*) can be changed, and changes must only be made with the appropriate **sysadmsh** selection.

The */etc/passwd* file was expanded into an adjunct Protected Password database, which stores the encrypted password and other security parameters about each user. Other security databases also contain usernames and the system does not function correctly if the databases are inconsistent. Using **sysadmsh** ensures that all the databases remain consistent.

The appropriate **sysadmsh** selections for changing fields in */etc/passwd* are described below. Also described are instances where an editor can be used, but this action is discouraged because careless editing can lead to corruption of the files.

5.2.1 Numerical Group ID

Use the Accounts→User→Examine:Identity selection of **sysadmsh**; the "Login group" field refers to the group name corresponding to the current numerical group ID of the user. Modify the group as required.

An editor can be used carefully in single-user mode without side-effects.

5.2.2 Comments

Use the Accounts→User→Examine:Identity selection of **sysadmsh**; the "Comment" field refers to the comment field of */etc/passwd* for the user. Modify the comment as required.

An editor can be used carefully in single-user mode without side-effects.

5.2.3 Home Directory

Use the Accounts→User→Examine:Identity selection of **sysadmsh**; the "Path" field refers to the home directory field of */etc/passwd* for the user. The Home directory radio buttons provide a number of options for changing a user's home directory. The Edit radio button is the simplest, allowing a new home directory to be entered in the "Path" field. See "Adding a User" in the "Administering User Accounts" chapter of the *System Administrator's Guide* for a full description of each radio button and how to set up shared home directories.

If an editor is used, errors can prevent a user from logging in. If the path leading to the home directory is not absolute, not searchable by the user or contains colons or \n or the directory itself does not exist, then the user cannot login.

5.2.4 Login Shell

Use the Accounts→User→Examine:Identity selection of **sysadmsh**; the "Login shell" field refers to the login shell field of */etc/passwd* for the user. Modify the shell as required. Pressing ⟨F3⟩ provides a point-and-pick list of default shells. A default shell may also be selected by entering the name of the shell without any preceding path. (The shell is, however, stored in */etc/passwd* with an absolute path.) If a default shell is chosen, then any startup files required by the shell (such as *.profile* for Bourne shell) are created for the user in his home directory.

If a default shell is not entered in the "Path" field, then **sysadmsh** checks that the shell entered has an absolute path, the path is searchable, and the shell is executable by the user. If no shell is entered then the login shell field in */etc/passwd* is left blank and the Bourne shell is started for the user when he logs in.

If an editor is used to specify a login shell for a user that would not pass the **sysadmsh** checks described above, then the user cannot log in.

5.2.5 Username

Changing a username is not supported. There is no **sysadmsh** option to allow the username to be changed.

If an editor is used to change an existing user name to a new name, the user cannot log in. Any **crontab**(C) jobs also fail.

5.2.6 Numerical User ID

Changing the numerical user ID is not supported. There is no **sysadmsh** selection to allow the numerical user ID to be changed.

If an editor is used to change an existing user ID the same problems occur as when editing the username.

5.2.7 Entries in /etc/group

A new entry in */etc/group* can be made with either the Accounts→User→Create selection when adding a new user, or the Accounts→User→Examine:Identity selection when changing an existing user's parameters. In both selections, when a group name is entered in either the "Login group" or "Groups" field that is not already in */etc/group*, a pop-up form appears requesting a numerical ID for the group. The range of valid group IDs is declared in */etc/default/authsh*, and the pop-up form provides a default of one greater than the highest already-used group ID or the minimum valid value, whichever is the greater. If you specify your own value then it must lie within the range declared in */etc/default/authsh* and must not already be assigned to an existing group.

An editor can be used carefully in single-user mode to create entries in */etc/group* without side effects.

5.2.8 Group Membership

Use the Accounts→User→Create selection to add a newly created user to a group, or Accounts→User→Examine:Identity to add an existing user to a group. The "Groups" field refers to the group membership in */etc/passwd*. Users should not be made members of groups corresponding to protected subsystems. Although this causes **sysadmsh** to grant the user the corresponding subsystem authorization, this is not the correct method and can cause problems that are difficult to diagnose.

An editor can be used carefully in single-user mode to change group membership.

5.2.9 Group names

Changing a group name is not supported. There is no **sysadmsh** option to allow the group name to be changed.

An editor can be used in single-user mode to change administrator-created group names. Group names corresponding to names of protected subsystems must not be changed. This would cause no one to have the subsystem authorization, and possibly stop the system from accepting any further logins.

5.2.10 Group IDs

Changing a group ID is not supported. There is no **sysadmsh** option to allow the group ID to be changed.

An editor can be used in single-user mode to change the IDs of user-created groups. IDs of names corresponding to names of protected subsystems must not be changed. The effects would be the same as changing the group name.

5.3 System Security and DOS-under-UNIX System

SCO UNIX System V/386 is designed to meet the requirements of the C2 level of "trust" as defined by the NCSC *Trusted Computer System Evaluation Criteria*. This means that all accesses of subjects (processes) to objects (files and devices) are subject to Discretionary Access Controls (DAC) and auditing.

DOS-under-UNIX system operations (such as those used in SCO VP/ix or Open Desktop) require special privileges that are accessed using the **v86init**( ) system call. This system call has the potential to bypass DAC and auditing in certain cases, although this does not affect normal use of the system because it can only be issued by a process either run directly from the *root* login or installed as **setuid** root.

When you install a product such as VP/ix, you install some **setuid** *root* commands that issue the **v86init**( ) system call. MS-DOS programs controlled by the emulation process, unlike those associated with SCO UNIX System V/386, have not been modified to satisfy the C2 requirements. Thus, a system with this software installed does not adhere to the C2 guidelines.

5.4 Use of Promains

Promains (protected domains, described in **promain**(M)) were first made available as part of the C2 security enhancements in Release 3.2.0. The promain mechanism, designed to restrict the execution of a potentially dangerous binary, is of limited use and will be removed in the next release. Promains are documented in the "Using a Trusted System" chapter of the *User's Guide*. At this time, it is still necessary for the system administrator to assign the **nopromain** kernel authorization for normal operation. This is assigned by default, but if the defaults are changed, you must be certain it is assigned. (**nopromain** means no promain restriction on program execution; the feature is used by taking away the authorization.)

5.5 Multiple Audit Directories

Two audit directories are defined by default in the root filesystem. When the root filesystem becomes full, the system switches to the second directory */tcb/auditttmp/audit2*. If you have multiple filesystems, it is recommended that you locate a second audit directory in a separate filesystem. For best results, do the following:

1. Delete the */tcb/auditttmp/audit2* directory using the **sysadmsh** System→Audit→Collection→ Directories→Delete selection.

2. Create a new directory on another filesystem using the **sysadmsh** System→Audit→Collection→ Directories→Create selection.

If this is done, the system will switch to the new directory on the other filesystem and prevent audit system shutdown due to the lack of space on the primary filesystem.

6. Using Your System

6.1 ct(C) Remote Terminal Program

The **ct**(C) program does not work in this release.

6.2 cu(C) Options

The following undocumented options are available for the **cu**(C) command:

~!*cmd...* runs *cmd* on the local system (via **sh -c**).

~$*cmd...* runs *cmd* locally and sends its output to the remote system.

~+*cmd...* runs *cmd* on the local system (via **sh -c**), with both standard input and standard output of *cmd* redirected to the remote system.

6.3 df(C) Options

The **df**(C) command has an undocumented **-i** option, which displays information on filesystem inode usage.

6.4 The Korn Shell: ksh(C)

ksh is the most recent version of the Korn Shell programming language, which is upward compatible with the Bourne Shell (with some exceptions; see "ksh and sh Incompatibilities") and has many added features. (Appendix B "Using the Korn Shell (ksh)" is a tutorial on using **ksh**.) **ksh** provides an enhanced programming environment in addition to the major command-entry features of the **csh**. With **ksh**, medium-sized programming tasks can be performed at shell level without a significant loss in performance. In addition, **sh** scripts can be run on **ksh** without modification. **ksh** also accepts 8-bit character sets transparently.

ksh provides the following features:

- Enhanced command re-entry capability. The **ksh** history function records commands entered at any shell level and stores them, up to a user-specified limit, even after you log off. This allows you to re-enter long commands with a few keystrokes - even commands entered yesterday. The history file allows for 8-bit characters in commands and supports essentially unlimited size histories.

- In-line editing. In **sh**, the only way to fix mistyped commands is to backspace or retype the line. **ksh** allows you to edit a command line using a choice of EMACS-TC or **vi**(C) functions. You can use the in-line editors to complete filenames as you type them. You may also use this editing feature when entering command lines from your history file.

- Extended I/O capabilities. **ksh** provides several I/O capabilities not available in **sh**, including the ability to:

  — specify a file descriptor for input

  — start up and run coprocesses

  — produce a prompt at the terminal before a read

  — format and interpret responses to a menu easily

  — echo lines exactly as output without escape processing

  — read and echo lines ending in \.

- Improved performance. **ksh** executes many scripts faster than the System V Bourne shell. A major reason for this is that many of the functions provided by **echo** and **expr** are built in.

- Integer arithmetic. A built-in command in **ksh** allows you to do integer arithmetic in any base from 2 to 36. Almost the complete set of C language operators are available. Further, variables in arithmetic expressions may include one-dimensional arrays. Arithmetic expressions can be used to form arguments to commands.

- Shell functions and aliases. Two mechanisms, functions and aliases, can be used to assign a user-selected identifier to an existing command or shell script. Functions allow local variables and provide scoping for exception handling. Functions can be searched for and loaded on first reference as with scripts.

- Substring capabilities. **ksh** allows you to create a substring of any given string directly by stripping off leading or trailing substrings during parameter substitution. You can also specify attributes, such as upper- and lowercase, field width, and justification, to shell variables.

- More pattern matching capabilities. **ksh** allows you to specify regular expressions for file and string matches.

- Improved debugging. **ksh** can generate line numbers on execution traces. Also, I/O redirections are now traced. There is a DEBUG trap that gets evaluated after each command so that errors can be localized.

- Job control. This feature allows you to stop and restart programs, and to move programs between the foreground and the background.

6.4.1 Job Control and Programs That Ignore It

If you run a program with job control set, there is a possibility it may not work properly. This is because the underlying code might not obey the job control specification (for example, the command **stty icanon &** would fail). If this occurs, simply disengage job control with the following command:

```
set +m
```

The program should then execute normally.

6.4.2 Job Control and Screen Redraw Problems

Because most programs do not take job control into consideration, the screen is not redrawn when the program is brought back into the foreground. When a screen-based application (one using the **curses** library) is placed in the background, bringing it into the foreground could result in a blank screen. When the screen is redrawn, the application should function properly.

6.4.3 ksh and shutdown(ADM)

If you attempt to run **shutdown su** from within **ksh**, the system will hang (freeze). The workaround is to start a Bourne shell first and then run the command.

6.4.4 ksh and sh Incompatibilities

The following is a list of known incompatibilities between **ksh** and **sh**:

- If an environment parameter is modified by **ksh**, the new value is passed to the children. In **sh**, you must export the parameter for this to happen.

- **time** is a reserved word in **ksh**. Thus **time a | b** times the pipeline in **ksh** while only **a** is timed with **sh**. Unlike **sh**, you can also time built-in commands and functions with **ksh**.

- **select** and **function** are reserved words in **ksh**.

- Parameter assignments only have scope for the command or function they precede in **ksh**. Only a subset of built-in commands in **ksh** treat parameter assignments globally. In **sh**, all built-in commands and functions treat parameter assignments as globals.

- The output of some built-in commands and error messages is different in a few cases; for example, **times** produces two lines of output in **ksh** and only one line of output in **sh**.

- While loops with redirection are not executed in a separate process in **ksh**, so assignments made within loops remain in effect after the loop completes.

- The semantics of functions are somewhat different. **ksh** can have local variables and allows recursive functions. Errors in functions abort the function but not the script that they are in. The parameter **$0** is the name of the function in **ksh**.

- The name space for functions and variables is separate in **ksh**. In **sh** they share the same space. The built-in **unset** requires a **-f** flag to unset a function in **ksh**.

- Words that begin with ~ may be expanded in **ksh**. **sh** does not have this feature.

- The character ^ is not special in **ksh**. In **sh** it is an archaic synonym for |.

- Whenever a command is surrounded by ((and)), **ksh** assumes that an arithmetic expression follows. In **sh** this means a subshell inside a subshell.

- Non-blank contiguous IFS delimiters generate a null input argument. Therefore, you can use IFS=: and correctly read the */etc/passwd* file even when fields are omitted. In **sh**, multiple delimiters count as a single delimiter.

- Arithmetic test comparison operators (**-eq**, **-lt**, ...) allow any arithmetic expressions. **sh** allows only constants. If you say **test x -eq 0** in **sh**, which is meaningless, it returns true, but in **ksh** it depends on the value of the variable **x**. If there is no variable **x**, then **ksh** produces an error message.

- The environment handed down to a program is not sorted in **ksh**. (Sorted environments are a quirk of **sh**. This should not be done because any user program can provide an environment list which does not have to be sorted.)

- The expansion of **$@** with no arguments produces the null string in the Bourne shell and produces nothing with **ksh** when there are no arguments. Setting **--** with no arguments unsets the positional parameter list in **ksh.** Thus, scripts that use set **--** **$@** when there are no positional parameters do not break.

- **ksh** accepts options of the form **-x -v** as well as **-xv** both for invocation and for the **set** built-in. The Bourne shell only allows one option parameter.

- **ksh** does not allow unbalanced quotes with any script. If the end of file is reached before a balancing quote in **sh,** it quietly inserts the balancing quote. **ksh** behaves like **sh** for **eval** statements.

- Failures of any built-in command cause a script to abort in **sh. ksh** scripts will only abort on errors in certain documented built-ins. In this respect **ksh** treats most built-in commands semantically the same as non-built-in commands.

- The sequence **$(** is special in **ksh.** In **sh** the sequence is illegal unless quoted. **"$("** must be preceded by a **** in **ksh** to remove its special meaning.

- The built-in command **exec** when used without arguments (for I/O redirection), closes on exec each file unit greater than 2.

6.5 maildelivery(F): User-Level MMDF Features

Users can create a **maildelivery**(F) file in their home directory to customize delivery of their mail. See the **maildelivery**(F) manual page in your *Replacement Manual Pages.*

6.6 mail(C) and the chron Option

If you have the **chron** option set in **mail** (messages displayed in chronological order), the **h** (headers) command does not display the last message in the mailbox. The message can be accessed by its message number or by the **$** (last message) designation.

6.7 Encryption Software Availability

The **crypt**(C) command and libraries are not distributed with the SCO UNIX System V/386 or Development System. Although the documentation is included, the distribution of encryption software to sites outside the United States is prohibited by the United States government. If you require the **crypt**(C) utility and associated **crypt**(S) libraries, and you are located within the United States, contact the support center listed on the support information card included with the software.

6.8 quot(C)

The **quot**(C) command does not work for users at this time. It can only be run by the super user.

6.9 rcvtrip(C) Mail Command

The **rcvtrip**(C) command can be used to notify other users when you are not available to answer your mail. See the **rcvtrip**(C) manual page in your *User's Reference Guide*.

6.10 su(C) Documentation

The manual page for **su**(C) does not state that a user name must be supplied on the command line when the **-c** option is used.

6.11 tic(C) Documentation

The **tic**(C) manual page includes a **-p** option, which is not valid. In addition, the undocumented **-c** option is used to check *file* for errors. Errors in **use=** links are not detected.

6.12 w(C) and uptime(C)

The **w**(C) and **uptime**(C) utilities now display load average information as documented.

7. Using the System Console

This section discusses the configuration and use of console displays.

7.1 Console Display Problems When Booting

When booting up, if the console screen blanks, the cursor is gone, or the display is garbled, you may have an incompatible video card. Check "Video Adapters and Monitors" in Appendix A of these notes for a list of supported video cards and monitors. If your video card is not listed, but is identical to (compatible with) one listed in Appendix A, it should work. If the card does not work, check the card hardware documentation. See if there are ways to configure the switch settings so the card is in an IBM-compatible emulation mode, and that it is addressing the kind of monitor attached. In particular, disable "autoswitch" modes. If changing the switch settings fails, then your monitor card is incompatible and must be replaced with a compatible card.

7.2 Console Selection When Booting

You can select the system console at boot time by entering

systty=cn

or place the keyword **systty=cn** in the file */etc/default/boot*.

7.3 Reducing the Number of Multiscreens

The system is configured with 12 **multiscreen**(M) console screens by default. Although this does not significantly affect performance, you can reduce the number of screens if desired. To configure the system for fewer screens, do the following:

1. Log in as *root* and enter the **sysadmsh**(ADM).

2. Use the System→Configure→Kernel→Parameters selection and select category 6, Multiscreens.

3. Skip the parameters displayed by pressing ⟨Return⟩ until you reach the NSCRN parameter. Enter the value corresponding to the number of screens you wish to configure.

4. Calculate the amount of screen memory in Kbytes (controlled by the SCRNMEM parameter) set aside for screens as follows:

   For 25-line displays:
   SCRNMEM = 10K + 4K * NSCRN

   For 43-line displays:
   SCRNMEM = 10K + 8K * NSCRN

5. Change the value of SCRNMEM to match the value you calculated in the previous step.

6. Exit the **configure**(ADM) menu by entering **q** and pressing ⟨Return⟩.

7. Use the System→Configure→Kernel→Rebuild selection to build the new kernel with the revised screen values. (The Rebuild should be visible when you exit the previous selection.) Follow the prompts. Exit **sysadmsh** when the process is complete.

8. You must then disable the unused **multiscreen** ttys. For example, if you reconfigured for 5 screens after having 12, you would enter the following command at the system prompt to disable ttys 6 through 12:

   disable tty06 tty07 tty08 tty09 tty10 tty11 tty12

9. You should now shut down the system and reboot from the new kernel. Enter the following command:

   shutdown -g0

   When the reboot message is displayed, press ⟨Return⟩ to restart the system.

7.4 Console Screen Protection and EGA Monitors

The console screen protection feature described in the "Using the System Console and Color Displays" does not work properly with EGA monitors; the screen will not blank after the prescribed number of seconds.

8. Using Printers

This section contains information about configuring and troubleshooting printers.

8.1 Slow Parallel Printers

If you have a parallel printer that prints abnormally slowly, check that your configuration matches the section "Installing a Printer" in the chapter "Using Printers" in the *System Administrator's Guide*.

If your printer is still slow, that is, taking about four seconds per line, your printer may be deselecting itself after receiving each line of text, and then reselecting itself for the next line. A kernel patch is provided to adapt your printer driver to this type of printer.

Slow printing has been observed on some models of Tandy® printers, but is not exclusively found on Tandy printers.

Note

The application of these patches can cause other classes of parallel printers to hang.

You can use this patch with either polled or standard lp devices.

Warning: The pound signs (#) and asterisks (*) are prompts from the system shell and from the patching program, */etc/_fst*. Do not type them in.

```
# cd /etc/conf/pack.d/pa
# cp Driver.o Driver.o.00
# /etc/_fst -w Driver.o
* patime+ad?xxx
_patime+0xad: 0x45f6 0x10f8 0xe74
* patime+ad?w 9090 9090 9090
_patime+0xad: 0x45f6=          0x9090
_patime+0xaf: 0x10f8=          0x9090
_patime+0xb1: 0xe74= 0x9090
* $q
#
```

8.2 lpsched(ADM) Options

Despite what is documented on the **lpsched**(ADM) manual page, the **lpsched** command does not accept any options. The options listed in the manual page and the usage message are not yet implemented.

8.3 Configuring a Network Printer

The **sysadmsh** option Printers→Configure→Add can set up the print spooler to print to remote printers that are directly connected to a modem. It does not configure the spooler to print to a printer on another machine accessed via network, such as TCP/IP, UUCP or Micnet. To do this, you must set up the printer as follows:

1. Choose Printers→Configure→Add from **sysadmsh**. Give the printer a name, and choose the existing interface script network. The Connection should be "Direct" and the Device should be "Hardwired." If you are adding a printer on a TCP/IP network, use */dev/null* for the device name. Otherwise, this should be left blank.

2. Create a file called */usr/spool/lp/remote*. This file should contain the following line:

   name:⟨TAB⟩*command*

   where *name* is the name of the printer on this machine and *command* is the command which will transfer the file to be printed to the remote system. For example, if the name you gave the printer is *remprint* and the name of the remote machine to which this printer is connected is *wally* and the name of the printer on machine *wally* is called *beaver* then the line would read:

   For UUCP:

   ```
   remprint: uux - wally!/usr/bin/lpr -dbeaver
   ```

   For Micnet:

   ```
   remprint: remote - wally /usr/bin/lpr -dbeaver
   ```

   For TCP/IP:

   ```
   remprint: rcmd wally lp -dbeaver
   ```

3. The file */usr/spool/lp/remote* should have owner and group **lp**, and be publicly readable. To set this, issue the following commands:

   chown lp /usr/spool/lp/remote
   chgrp lp /usr/spool/lp/remote
   chmod 444 /usr/spool/lp/remote

4. For a TCP/IP network, note that the user wishing to print will need permission to invoke the remote command on this machine. Again using TCP/IP as an example, this requires an entry for the remote machine in */etc/host.equiv* file for each machine on the network that can request a print job. In addition, each user must have an account on the remote print server machine. One limitation is that the *root* user cannot send a print request.

8.4 lp(C) -c Option

The **-c** option is obsolete because this feature is engaged by default.

9. Using Floppies and Tapes

This section discusses floppy and tape drives.

9.1 Configuring a SCSI Tape Drive

The system is configured for a non-SCSI cartridge tape drive by default. If you add a SCSI tape drive, you must remove the existing configuration first. When invoking **mkdev tape**, you must select the Remove a Tape Drive option before adding the SCSI tape drive. See the "Using Floppy Disks and Tape Drives" chapter of the *System Administrator's Guide*.

9.2 New Irwin Driver

This release includes an enhanced Irwin driver. Several commands were added to the **tape**(C) command and the new Irwin tape **ioctl**(S) commands are documented in **tape**(HW). These pages are located in your *Replacement Manual Pages*.

9.3 The tape load and unload Commands

The **load** and **unload** options of the **tape** command are used on certain drives to load and unload the tape. Drives that support **load** and **unload** include:

9-track drives
DAT drives
Wangtek 5525ES
Exabyte 8mm

With 9-track drives and the Wangtek 5525ES, a **tape load** command must be issued before any other tape commands. The Exabyte 8mm and most DAT drives do an automatic load when the tape cartridge is inserted.

On tape drives that do not support automatic loading and unloading of media, **tape load** and **tape unload** typically just turn the front panel light on and off.

9.4 Adding a SCSI Tape Drive

If you want to add a SCSI tape drive with a non-standard adapter, you must follow the instructions in the section of these notes entitled "Using a Non-Standard SCSI Adapter".

If you want to add a tape drive and are using an Adaptec SCSI card, or you have other SCSI devices already configured, follow the instructions below:

1. Log in as *root*.

2. Enter the command:

   mkdev tape

3. Select the Remove a Tape Drive option and remove the cartridge tape driver.

4. Next, select the Install a Tape Drive option and install the SCSI tape driver.

5. When asked whether you want to relink the kernel, respond **n** for no. The SCSI driver is already linked into the kernel.

If you attempt to make the SCSI drive the default tape drive after installing a cartridge tape drive, you must again remove the cartridge tape drive so that the proper device nodes can be created.

The **mknod**(C) command is used to add new device nodes. You must supply a major and a minor device number. To create a "no unload" device for your SCSI tape drive, you must use the major number of your normal tape device, and add 4 to the minor device number. Use the **ls -l** command to list the major and minor device numbers. For example, to create the "no unload" counterpart to a */dev/rct0* with a major number of 46 and a minor number of 0, follow these instructions:

1. Log in as *root*.
2. Enter the command:

   mknod */dev/nurct0* **c 46 4**

This creates a device called */dev/nurct0*. You must:

- Reference this device explicitly when you access your tape drive.
- Add the "no unload" device name to */etc/default/tar* or */usr/lib/sysadmin/schedule* (for scheduled backups) as desired.

Note that you do not need to use a "no unload" device if you are already using a "no rewind" device, as this device does not automatically unload the tape.

9.5 Creating Backups with Irwin and QIC-40 Drives

We recommend that tape backups be done in single-user mode for most efficient use. If you see an error message stating "cannot allocate buffer" or "not enough space" while using the tape drive, reboot your system, enter system maintenance (single-user) mode and run the backup again.

9.6 SCSI Tape Drives and ECC

Tape ECC (Error Correction Code) support is not available for SCSI tape drives at this time.

9.7 SCSI Tape Device Nodes on GT Floppy Distribution

The minor device numbers for the SCSI tape device (*/dev/rStp0*), and the no-rewind SCSI tape device (*/dev/nrStp0*), are reversed (GT floppy distributions only). Running **mkdev tape** and adding a SCSI tape drive creates the devices properly, but this can also be done manually by logging in as *root* and entering the following commands:

```
cd /dev
mv rStp0 temp
mv nrStr0 rStp0
mv temp nrStp0
```

10. Using a Mouse

This section describes features or problems that affect the use of the mouse.

10.1 Using mscreen(M) and usemouse(C)

If you are using **mscreen**(M) and **usemouse**(C) simultaneously, these two programs compete for the same pseudo-tty devices. To avoid potential conflicts, use the *.mscreenrc* file documented in the **mscreen** manual page. The *.mscreenrc* file allows you to assign a specific set of pseudo-tty devices to individual users. Each user using **mscreen** should have an *.mscreenrc* file in their home directory.

10.2 Installing a Keyboard Mouse

A keyboard mouse is a mouse that connects to a special port found on some new AT-type and EISA machines, and all Microchannel machines. When installing a mouse using **mkdev mouse**, choose the keyboard mouse option only if you have a machine and mouse capable of using this port. On some AT-type machines that do not have a keyboard mouse port, installing a keyboard mouse may cause the keyboard to lock up after booting. If this happens, boot the previous kernel (**unix.old**) and remove the keyboard mouse driver from the system using the remove option of **mkdev mouse**. This problem does not occur on systems with keyboard mouse ports.

11. Using MS-DOS and OS/2

This section concerns compatibility issues with MS-DOS and OS/2.

11.1 UNIX and MS-DOS Coexistence

The SCO UNIX System V/386 supports the coexistence of MS-DOS™ on the same hard disk (only MS-DOS 3.3 or earlier; 4.0 extended partitions are not supported). Some versions of MS-DOS have restrictions; for example, ITT DOS releases previous to 3.10 cannot share the disk with the UNIX partition or MS-DOS. For these releases, the UNIX partition must occupy the whole disk.

As another example, an ITT MS-DOS+ release 3.20 hard disk partition should not be made larger than 32 Mbytes. Activating a larger partition corrupts the MS-DOS+ ending cylinder. UNIX **dos**(C) commands may not work when accessing an ITT MS-DOS+ release 3.20 hard disk partition that is 32 Mbytes or larger.

Likewise, some versions of NCR MS-DOS only recognize 32-Mbyte partitions. You see the message "No operating system on fixed disk" when attempting to boot an MS-DOS partition larger than 32 Mbytes.

Whenever you use MS-DOS and the SCO UNIX System V/386 on the same disk, if in doubt, install MS-DOS first, then install the SCO UNIX System V/386.

11.2 MS-DOS Filesystem Support

MS-DOS (3.3 and earlier) floppy and hard disk filesystems are now mountable from the UNIX partition, and their files are accessible using standard file utilities. Support for this feature must be linked into the kernel. See the "Using DOS and OS/2" chapter of the *System Administrator's Guide* for details.

12. Using Networks

This section contains notes about networks.

12.1 UUCP and System Security

Installation of the UUCP system on your system violates the requirements for the C2 level of trust. This includes the **cu**(C) program. There are several reasons for this:

- No accountability. There is no pairing of an individual with a single UID on the system. A popular practice with UUCP accounts are anonymous logins without passwords, which pose an even greater threat.
- Potential risks from imported files. Unless you set the system up carefully, files can be placed on your system that pose a risk to security, as well as the possibility of undesired access to your files.
- No security standard exists for networks. At this time, there are no standards for a trusted network. The specifications for such a network are still in the design stage.

If you do not care about C2 guidelines, and if you set up your UUCP system with caution (using the facilities of the UUCP **permissions**(F) file), you can operate UUCP with minimal risk.

12.2 UUCP Address Conventions

You should never mix the INTERNET @ (at symbol) and the UUCP ! (exclamation point or bang) conventions when addressing mail. Either convention is acceptable, as in the following examples:

```
ripley@nostromo.COM
nostromo.COM!ripley
```

Do not try a construction like this:

```
forbin!colossus@cpo
```

12.3 UUCP Anonymous Login Accounts

If you have an anonymous UUCP login account **uucp** that you are using for anonymous logins, you should change this account to **nuucp**, which is the proper name for an anonymous UUCP login. Although use of anonymous logins is not recommended, no account should ever be without a password, and you should never use **uucp** as an anonymous login account.

UUCP login accounts are created with a default password expiration of 14 days. To alter this, you must use the **sysadmsh**(ADM) Accounts→User→Examine:Expiration selection to redefine this limit. For more information, see the "Administering User Accounts" in the *System Administrator's Guide*.

If you have difficulties with UUCP accounts being locked (messages like "dead account" are displayed), you can extend the number of login attempts by selecting Accounts→User→Examine:Logins. If the account was locked due to too many unsuccessful login attempts, the "Account Locked" field displays "Too many unsuccessful login attempts." You can clear this condition either by setting the maximum unsuccessful logins to a larger number (including infinite), or by selecting the Lock status option Clear all locks.

12.4 Third Party Modem Communications Programs and UUCP

The UUCP package (which includes **cu**(C)) must have ownership and exclusive write permission on ttys. Because the UUCP package sets permissions on tty devices used with modems, any changes made to accomodate other communications programs (such as **kermit**) will be altered by UUCP communications that access the same tty. For example, if tty1A is changed to include group write (664) permissions for use with another communications program, UUCP will restore the permissions to exclusive write (644). The solution is to change the binaries for other communications programs to be setuid *uucp*. To change the permissions on other programs, log in as *root* and enter the following commands (substituting the actual program name for *program*):

chown uucp *program*
chmod u+s *program*

13. System Configuration and Link Kit Notes

This section discusses issues relating to system parameters, drivers, and the Link Kit.

13.1 System Priority Level and Driver Performance

In SCO UNIX Release 3.2.0, the kernel routines dealing with clists and streams used system priority level 7 to protect their critical sections. However, the use of this extremely high level was found to impair network bandwidth. Under 3.2v2.0, the clist and stream priority levels now use their traditional value of 5. We have checked with a number of board and driver vendors, and do not know of any whose drivers interlock using spl7( ). A driver depending on spl7( ) interlocking could fail under the new scheme. If you are using a vendor-supplied driver and experience periodic hangs or otherwise unexplainable system panics when that driver is in use, you may want to patch your kernel so that interlocking is done at spl7( ). You can use the following script to change the priority level:

1. Shut down the system and bring it up in single-user mode. You can do this by invoking the **shutdown**(ADM) command and entering the *root* password when you are prompted to press ⟨CTL⟩**d** or enter the *root* password.

2. Enter the following commands:
   cp /unix /unix.spl7
   /etc/_fst -w /unix

3. The system prompt then changes from a crosshatch (#) to an asterisk (*). Enter the following command:
   spltty+1?d

4. The system should respond with the following:
   ```
   spltty+0x1:7
   ```
   If the value at the right is anything other than 7, do not continue this procedure. Enter **$q** to exit */etc/_fst*.

5. Next enter:

   spltty+1?w 5

6. The system should respond:

   ```
   spltty+0x1:  0x7= 0x5
   ```

7. Enter the following command to quit:

   $q

The change takes effect the next time you boot the system. If this does not solve your problem, or this procedure causes other problems worse than the ones you started with, make certain you are logged in as *root* and enter the following command:

mv /unix.spl7 /unix

This restores the system to its former state after you reboot the system. Note that if the procedure was successful and you relink a new kernel at a later date, you must repeat this procedure.

13.2 10-bit I/O Addressing Check

The operating system detects machines that have only 10-bit I/O addressing. This message is displayed at boot time:

```
kernel: INFO: 10 bits of I/O address decoding
```

Such machines cannot use bus cards at addresses above 0x400. Card addresses must be below 0x400 on machines with only 10 bits of I/O address decoding.

13.3 NDISK Parameter Description

The description of NDISK, the tunable kernel parameter, found on pages 18 through 24 of the "Tuning System Performance" chapter of the *System Administrator's Guide* is incorrect. The NDISK parameter is not set at boot time. The default value of NDISK is 4 and must be incremented if more than 4 disks are added to the system.

14. Internationalization

This section deals with internationalization issues.

14.1 Internationalized Utilities

The **vi**(C), **sh**(C), and **mail**(C) utilities were modified to handle 8-bit characters.

14.2 Documented Features Not Present in This Release

14.2.1 file(C)

The ability to identify 8-bit text is not yet available.

14.2.2 mail(C)

mail has only been partially internationalized; it does not accept 8-bit characters in user names.

15. Compatibility and Conformance Notes

This section contains information about compatibility with other UNIX Operating System and XENIX releases, including conformance to industry standards.

15.1 Security Standards Conformance

The security features present in this release are extensions to System V UNIX and are designed to meet the requirements of the "Class C2 Controlled Access Protection" rating as specified in the following security standards:

- DhH 5200.28-STD "Department of Defense Trusted Computer System Evaluation Criteria" (also known as the "Orange Book").

- TD-85-02 "Department of the Treasury Handbook for Automated Information Systems Security and Risk Management, April 1987".

Furthermore, password management facilities are present that are designed to meet the following standards:

- CSC-STD -002-85 "Department of Defense Password Management Guideline" (also known as the "Green Book").
- FIPS PUB 112 "Password Usage" (a superset of the requirements described in the "Green Book").

15.2 XENIX System V Compatibility

SCO UNIX System V/386 provides full binary and source code compatibility with applications developed for XENIX System V/386, XENIX System V/286, UNIX System V/386, and UNIX System V/286. This support does not extend to device drivers, however. The following list describes the level of XENIX System V support:

- Source code written for XENIX System V/386 can be compiled and linked on UNIX System V/386 without having to modify the source code.
- Binary applications developed for XENIX System V/386 (Release 2.2.0 and later), XENIX System V/286 (Release 2.0 and later), and XENIX System V/8086 (Release 2.0 and later) can be run on UNIX System V/386 without having to recompile the applications.
- The structure of the UNIX System V/386 filesystem allows both XENIX System V and UNIX System V binary applications to be executed. It also supports the mounting of XENIX and UNIX removable filesystems.
- Support for XENIX system call extensions enables programs to run as they did under the XENIX system.
- All device driver support routines available under XENIX System V/386 are available in UNIX System V/386.

15.2.1 Differing System Calls

The following XENIX system calls are supported in SCO UNIX System V/386 but function differently:

ptrace()
In SCO UNIX System V/386 the **ptrace()** system call is not supported for XENIX system binaries. XENIX system binaries that rely on **ptrace** do not work on SCO UNIX System V/386. To make their code run on SCO UNIX System V/386, you must modify their XENIX system code to use the UNIX System version of **ptrace**.

ulimit()
XENIX System V/386 binaries that call **ulimit()** with the cmd argument set to 2 cannot increase their limit beyond the maximum number of blocks that are representable in a 512-byte block filesystem. This restriction is not enforced when the source is compiled on SCO UNIX System V/386.

uname()
The utsname structure returned from **uname()** is a different size, depending on whether you compile on XENIX System V/386 or on Release 3.2. On XENIX System V/386, there are extra fields at the end of the structure.

15.2.2 XENIX-286 Application Execution

This release of the UNIX System contains a XENIX System V/286 utility that allows XENIX System V/286 (Release 2.3.x) programs to run transparently on the Intel 80386 processor under UNIX System V/386 Release 3.2.

For more information about the XENIX-286 feature, see the **x286emul**(C) manual page in the *User's Reference.*

AT&T and XENIX conventions for device names are completely supported. The directories and naming schemes are discussed in the "UNIX Directories and Special Device Files" chapter of the *System Administrator's Guide.*

15.3 AT&T SVID Conformance

The AT&T System V Interface Definition (SVID) is a three-volume set of books published by AT&T (select code 307-127). SCO UNIX System V/386 was tested against the AT&T SVVS3 (System V Verification Suite) tests, and the conformance was verified.

15.4 POSIX P1003 Conformance

The referenced standard is published by the IEEE. SCO UNIX System V/386 conforms to the POSIX P1003.1 Operating System Interface Specification with only minor exceptions.

15.5 FIPS PUB 151-1

SCO UNIX System V/386 substantially conforms to the FIPS PUB 151-1 requirements, and future releases will be in full conformance.

15.6 The ISO 8859 Character Set

The operating system can handle all programs using the ISO 8859 character sets. It provides full 8-bit support and is configurable for 7- and 8-bit peripherals with different character sets, including mapping files and conversion tools.

15.7 X/Open CAE Conformance

The X/Open specification is defined in the X/Open Common Applications Environment (CAE) *X/Open Portability Guide*, published by X/Open, Ltd. This specification includes systems and applications software, including languages and database systems. SCO offers a full X/Open CAE-compatible product line. The operating system portion is defined by the *X/Open Portability Guide* (XPG).

In addition, the XPG specifies 8-bit libraries and 8-bit versions of commands for the international market. These features were added to SCO UNIX System V/386.

16. Programmers Notes

This section contains notes to device writers, utility and shell script writers.

16.1 Notes to Device Driver Writers

If you are writing a device driver which must run on both industry standard (AT) and microchannel (MC) architectures, but will operate differently depending on architecture type, you must use the **extern unsigned long arch** to determine which type of architecture the kernel is running on. The **AT** bit will be set if the kernel is running on industry standard architecture. The **MC** bit will be set if the kernel is running on microchannel architecture. Refer to the include file */usr/include/sys/arch.h* for the declaration of **arch** and the definition of the **AT** and **MC** bits. Note that you must compile your source using **-DARCH=AT+MC** on the **cc** command line, or with **#define ARCH AT+MC** in the source file.

If you are writing a device driver for microchannel architecture, which needs to refer to the configuration POS registers, call **cdbinit()** and then use the byte array **mcapos [8] [8]**. Refer to the include file */usr/include/sys/cram.h* for the declaration of **mcapos** and it's associated definitions.

Note that **arch** and **mcapos** were not defined in SCO UNIX System V/386 Release 3.2.0, nor in SCO Open Desktop Release 1.0. **arch** is defined and set by both SCO UNIX System V/386 Release 3.2 Version 2.0 GT and SCO UNIX System V/386 Release 3.2 Version 2.0 MC. Additional architecture bits may be supported in later releases. **mcapos** is not defined in SCO UNIX System V/386 Release 3.2 Version 2.0 GT but is defined and set in SCO UNIX System V/386 Release 3.2 Version 2.0 MC and later releases.

16.2 Notes to Utility and Shell Script Writers

SCO UNIX System V/386 Release 3.2 Version 2.0 MC includes the device */dev/mcapos* (read only character device, major device number 5, minor device number 2). The 64 bytes of MCA POS register configuration information can be read from this device. Refer to the include file */usr/include/sys/cram.h* for a declaration of the kernel byte array **mcapos [8] [8]** from which these are read, and it's associated definitions.

Note that */dev/mcapos* was not supported under SCO UNIX System V/386 Release 3.2.0 nor in SCO Open Desktop Release 1.0. A device of major 5 minor 2 will not give correct information in those releases. */dev/mcapos* is not defined in SCO UNIX System V/386 Release 3.2 Version 2.0 GT. Therefore, absence of */dev/mcapos* may be taken to indicate that the machine is not based on microchannel architecture. Reading from device major 5 minor 2 under SCO UNIX System V/386 Release 3.2 v2.0 GT will return an error, setting **errno** to **ENXIO**, unless your computer is based on microchannel architecture.

SCO UNIX System V/386 Release 3.2 Version 2.0 MC also includes the **slot** utility */etc/slot*. **slot** reads */dev/mcapos* and the text file */etc/default/slot*, to display the MCA POS register configuration information. See the manual page **slot**(C) in your *User's Reference Guide* for full details of this command.

Appendix A
SCO UNIX System V/386
Release 3.2 Version 2.0
Compatible Hardware

A.1 Compatible Hardware

SCO UNIX System V/386 is available for many configurations of personal computer hardware, for both Industry Standard, Extended Industry Standard (EISA), and Microchannel architectures. (EISA and Industry Standard are considered equivalent except where noted.) Your hardware configuration must have the original settings and boards before you install the Operating System. If you have added any boards, make sure that all switches are set as recommended in the manufacturer's hardware manual for that board.

This appendix is divided into several sections. The first few sections list the supported machines by microprocessor and architecture type. "General Compatibility Guidelines" contains general information on hardware compatibility that apply to both architectures.

The rest of the appendix is divided into the Standard Architecture and Microchannel Architecture Hardware Notes. Each covers specific configuration details, including charts of compatible peripherals, serial cards, video adapters, monitors, hard disks, and controllers. System parameters necessary for these devices also appear in this section. These guidelines must be followed to ensure proper system performance.

Note

The specific hardware that is listed in these *Release Notes* has been used with SCO UNIX System V/386. However, because the manufacturers of compatible machines or add-on peripherals may change configuration, functionality, or firmware at any time, no guarantee is implied. Please write to us with detailed hardware information for possible inclusion on our lists.

A.2 Using This Appendix

To find a listing of compatible hardware for your machine, you must know your machine's classification. You must know the processor your machine uses and whether it uses Standard or Microchannel bus architecture. (Microchannel architecture refers to computers that are compatible with the IBM PS/2 computers.) You should also know if there are enhancements to your system, such as a non-ST506 AT disk controller. To find the available classifications of machines, read the next section, "Supported Hardware Configurations." There is a specific section of compatible hardware in this appendix for each version of SCO UNIX System V/386. There is also a section called "General Compatibility Guidelines" that discusses general compatibility issues with SCO UNIX System V/386. Read through this section before you install any extra hardware on your system.

A.3 Supported Hardware Configurations

SCO UNIX System V/386 works on standard 386-based personal computers, including those with support for Adaptec AHA-154x and Western Digital WD7000 SCSI Host Adapters and SMS/OMTI 862x Direct ESDI disk controller. Be sure to consult the "Disk Controllers and Host Adapters" section under "General Compatibility Guidelines" and the "SCSI Guidelines" section under each architecture for important information on specific SCSI adapters.

Some computers arrive with the hard disk only partially formatted. If you have such a machine, use the correct low-level or hard format procedure as described in the manual for your hard disk controller before installing SCO UNIX System V/386.

Some computers require specific switch settings to run SCO UNIX System V/386. If your computer does not run SCO UNIX System V/386 with the settings as shipped, contact your computer hardware representative for the proper settings.

A.4 Standard Architecture 386-Based Machines

We have used the following machines under SCO UNIX System V/386:

ALR FlexCache 20/386
ALR FlexCache 25/386
ALR FlexCache 33/386 Model 150
ALR Microflex 7000
ALR PowerCache 486e
AST Premium 386
AST Premium 386/33
AST Premium 486/25T
Laguna Systems PDQ386
Mitac 386
ACER 1100SX
ACER 1100/20†
ACER 1100/33‡
ACER System 32/20
Compaq 386
Compaq 386 33 Mhz
Compaq Portable III 386
Compaq DESKPRO 386/20
Compaq DESKPRO 386/25
Compuadd 386
CSS 386
DECstation 320
EasyData 386 model 333
Epson Equity 386SX
Epson Equity 386/25††
Epson Equity 386/20‡‡
Everex 386/20
Goupil uniprocessor 25MHz Tower
GRiDCase 1530
Hertz 386/25
Hewlett-Packard Vectra QS/16S
Hewlett-Packard Vectra QS/20
Hewlett-Packard Vectra RS/16 PC
Hewlett-Packard Vectra RS/20 PC

Hewlett-Packard Vectra RS/20C
Hewlett-Packard Vectra RS/25C
Intel 302
ITT 386
Legacy 386/33
Microflex 3300
Mitsuba 386
Mitsubishi PC-386
NCR 3386
NCR 316
NCR 316SX
NEC BusinessMate 386
NEC BusinessMate 386/33
NEC PowerMate 386
NEC Powermate 386/25Mhz
NEC 386 20 Mhz
Noble 386 from PC Discount
Nokia Alfaskop System 10 m52
Nokia Alfaskop System 10 m54/m55
Olivetti M380
Olivetti M380 XP5
Olivetti M380 XP7*
Olivetti M380 XP9*
PC Craft PCC 2400 386
Siemens Data Systems Model WX200
Tandy 4000
Tatung Force 386X
Tatung TCS-8000 386
Tatung TCS-8600 386
Televideo 386/25
Texas Instruments System 1300
Toshiba T3100SX
Toshiba T3200SX
Toshiba T5100 386
UNISYS PW2 Series 800/20
Victor 386 25Mhz
Wang PC 380

Wyse 386
Zenith SuperSport 386sx
Zenith Z-386/16
Zenith Z-386/25
Zenith Z-386/33
Zenith TurbosPORT
Zenith 386/20 model ZVB 3524ED

* On the XP7 and XP9, SCO UNIX System V/386 does not boot from the floppy drive if the Floppy Disk Access Speed Option is set in slow mode. The slow mode option is provided for the MS-DOS environment, and is not intended to be used in the SCO UNIX System V/386 environment. To boot your system from the floppy disk, ensure that this option is set to fast mode.

† The ACER System 15 Model 10 is equivalent to the ACER 1100/20.

‡ The ACER System 15 Model 20 is equivalent to the ACER 1100/33.

†† The Epson PC AX3/25 is equivalent to the Epson Equity 386/25.

‡‡ The Epson PC AX3 is equivalent to the Epson Equity 386/20.

The Intel Inboard in an IBM PC/AT or Intel-supported 286AT compatible is also supported.

The Orchid JET 386 in an IBM PC/AT or Orchid-supported 286AT compatible has also been reported to run SCO UNIX System V/386.

The AOX Master 386 in an IBM PC/AT or AOX-supported 286AT compatible has also been reported to run SCO UNIX System V/386.

These machines have been reported to run SCO UNIX System V/386, but we have not tested them:

Cheetah CAT-386
Corvus 331
Arnet Multiuser 386
NCR PC916

Note that the personality card supplied by NCR must be configured so that the Video Adapter auto-switch feature is disabled.

A.5 Microchannel Architecture 386-Based Machines

We have used the following machines under SCO UNIX System V/386:

- Apricot Qi 300 (requires additional memory)
- Apricot Qi 600
- Apricot Qi 900
- Apricot Vx FtServer range
- ALR 486MC Model 150
- IBM PS/2 Models
  - 55-031
  - 55-061
  - 65 SX-061
  - 65 SX-121
  - 70-E61
  - 70-121
  - 70-A21
  - 80-041
  - 80-071
  - 80-111
  - 80-311
  - 80-X21
  - 80-M21
  - 80-A21
  - 80-A31
- Tandy 5000MC
- Olivetti P800
- Olivetti P500

A.6 General Compatibility Guidelines

This section describes hardware that can be used with SCO UNIX System V/386. It contains hardware information that is generic to Industry Standard and Microchannel Architectures. There are many other devices that you can use, which require additional vendor supplied software, that are available from independent hardware vendors. Call your dealer or sales representative and ask for the third-party IHV/ISV catalogue. If your computer is listed as a supported machine in this appendix, it should run SCO UNIX System V/386 without adding any hardware or changing any jumper or switch settings, unless otherwise stated.

Note

Supported machines are not always supplied with video cards by the same manufacturer. Check the video card for compatibility.

A.6.1 Math Chips

Your personal computer may include the 80387 math coprocessor, which is automatically detected and supported by SCO UNIX System V/386. These coprocessors improve floating point efficiency. The 80486 CPU includes an on-chip coprocessor that is also recognized and used as an 80387.

Use math coprocessors matching your machine's CPU speed. Follow the manufacturer's recommendations.

At boot time, SCO UNIX System V/386 announces the presence of a math coprocessor with the message:

```
%fpu  - 35 - TYPE=80387
```

Please note that switches on the main system board must be set properly to enable 80387 interrupts and/or your system must be set up with the manufacturer's setup disk to expect the chip. Ensure that the system diagnostics recognize the coprocessor presence and check your hardware manual for the proper switch settings.

Please note that on some motherboards, the operating system incorrectly recognizes the presence of an 80387 coprocessor even if the chip is not installed. This problem is prevalent on machines that use the Intel motherboard; if your computer incorrectly recognizes the presence of an 80387 chip, make sure that blocks E48 and E49 are not connected with a jumper connection.

Some 80387 exceptions are masked. Refer to the manual page for **80387**(HW).

A.6.1.1 Weitek Coprocessors

The Weitek 1167 and 3167 numeric coprocessors are also supported. This support extends only to runtime; there is no current development support for creating binaries that take advantage of the 1167 and 3167.

In order for the Weitek chip to be recognized by the system, one file needs to be edited in the following way:

1. In the file */etc/conf/sdevice.d/weitek*, there should be a line similar to the following:

   ```
   weitek N 1 0 0 ...
   ```

   Change the "N" (for no) to a "Y" (for yes).

2. Issue the following commands to relink the kernel:

   cd /etc/conf/cf.d
   ./link_unix

   Answer "yes" when asked if this kernel should boot by default; answer "yes" again to rebuild the UNIX kernel environment.

3. Reboot the machine; the Weitek chip will be recognized.

A.6.2 Memory Cards

In general, most memory cards work with SCO UNIX System V/386. If you experience "panic: parity" errors it is often because of low-quality memory chips or cards. This problem is especially prevalent with the 32-bit static RAM chips used in older 386 machines.

With memory cards, check the switch settings on both the card and motherboard. Refer to the hardware manuals for your computer and for the memory card to find the correct switch settings. SCO UNIX System V/386 supports up to 64 Mbytes of main memory.

Use of 32-Bit Memory

It is *strongly* recommended that you use 32-bit memory from your machine manufacturer. 16-bit memory is much slower, and will degrade overall machine performance.

Note

Certain manufacturers reserve the upper 384K of the first megabyte for MS-DOS. On some machines, this "shadow" RAM cannot be accessed by SCO UNIX System V/386. You must install additional memory to run SCO UNIX System V/386.

If you see this message:

```
panic: memory failure -- parity error
```

some part of your hardware is sending a "non-maskable interrupt" (a signal sent by the hardware that halts the operating system). You should run your system's hardware diagnostics tests if available. In addition, physically re-seat your memory cards and chips, and check for bent pins, and so on. If these measures fail to correct the problem, or you do not feel comfortable in checking your hardware yourself, seek assistance from professional hardware experts. One frequent cause of the problem is memory chips that are slower than factory-recommended chips.

A.6.3 Multi-Function Cards

The serial ports on many multi-function cards behave as expected if COM1 and COM2 are fully compatible with the standard specifications for these serial ports. Memory, parallel ports, and other hardware usually function as expected.

A.6.4 Serial I/O Boards

This section describes the conditions and results of using various serial I/O boards with SCO UNIX System V/386. Standard single port serial I/O boards function as expected if COM1 and COM2 are fully compatible with the standard specifications for these serial ports.

To configure the system for the serial board you are installing, you must run the **mkdev serial** command. See the "Adding Multiport Cards, Memory, and Other Bus Cards" chapter of the *System Administrator's Guide* for more information on **mkdev serial**.

Note

If you are installing a "smart" multiport card (such as the Arnet Smartport card), do not use **mkdev serial** to configure your card. The manufacturer provides configuration software that is **custom**(ADM)-installable. Consult your hardware documentation for additional information.

Each multiport serial I/O board is unique; SCO UNIX System V/386 has a special driver code for each card listed. Only those with status poll registers can work with the high performance driver scheme chosen, and new boards require additional driver support.

Refer to the **serial**(HW) and **mkdev**(ADM) manual pages for more information on compatible serial I/O cards and on adding and enabling serial lines.

Serial I/O Chip Notes

Some computers or add-on serial I/O cards use the 8250a serial I/O chip. (16450 chips are strongly recommended for 386 machines.) Some revisions of this chip do not handle interrupts properly. MS-DOS does not use interrupts, so the use of this chip with MS-DOS causes no problems. SCO UNIX System V/386 makes use of interrupts, as it is a multitasking operating system.

The problem with the serial I/O chip shows up when using **uucp**(C) or **cu**(C). Indications that your computer contains a bad revision 8250a chip are that **uucp**(C) may lose characters constantly and generate unkillable **uucico** processes, and that **cu**(C) at high baud rates stops executing and does not exit.

The problem rarely shows up when using the serial port with a terminal. It is associated with high-speed serial input. If you want to use **uucp**(C) or **cu**(C) and your computer has one of these chips, we recommend you replace the 8250a chip with an 8250b serial I/O chip or use a multi-function card containing a serial port and configure it as COM1 or COM2. Disable the built-in serial port or avoid high-speed input on that port.

All COM1 boards should be strapped at Interrupt Vector 4. All COM2 boards should be strapped at Interrupt Vector 3. Check your serial card hardware manual or call the hardware manufacturer for the switch settings that implement these addresses.

Note

SCO uniPATH SNA-3270 uses Interrupt Vector 3, which can interfere with the use of a serial card on COM2.

A.6.5 Tape Drives and Controllers

Note that the tape drives described in these notes are sometimes sold under other brand names. Only one cartridge tape subsystem is supported per computer. (It is possible to have one cartridge tape drive plus a minicartridge or QIC-40 drive configured on a system.) This limit does not apply to SCSI tape devices. (Up to four SCSI tape drives are supported on a SCSI bus. SCSI tape drives can be configured at any ID on either Host Adapter.) QIC-24 format is supported on the full size cartridges. Tape support is raw (character) only, no block device. A no-rewind device exists for writing multiple tape files on a single tape. A tape utility, **tape**(C), is provided for rewinding, erasing, format, retensioning, and so on. Digital audio tape (DAT) and 8mm (video) tape drives are also supported as SCSI devices.

To configure your system for a tape unit, run **mkdev tape**. If you do not choose specific values for the DMA channel, interrupt, and base address, the default values are used. (This information is found in "Tape Drive/Controller Combinations" of this appendix in the section that applies to your machine architecture.) Note that most addresses are specified in hexadecimal. If you do not use the default settings, watch for possible interrupt conflicts with other installed devices. The "Tape Drive/Controller Combinations" section indicates which interrupts may be in use on your system. You should not use interrupts 0, 1, or 6, as these are already in use regardless of additional devices. The **mkdev tape** menu indicates the Irwin units as "Mini-Cartridge," and all other units are indicated as "Cartridge," "QIC-40," and "SCSI" units. The Irwin units are not configurable. Refer to the "Using Floppy Disks and Tape Drives" chapter of the *System Administrator's Guide*.

A.6.6 Video Adapters and Monitors

Any video adapter that truely emulates a standard adapter should run under SCO UNIX System V/386.

SCO UNIX System V/386 supports the use of two video adapters and two monitors, one of type mono and the other of type color. Use of only a single VGA adapter is supported.

Adapters bundled with supported computers (monochrome or color monitor) work with SCO UNIX System V/386.

A.6.7 Add-On Hard Disks

Many hard disks, both standard and nonstandard, can be used by the SCO UNIX System V/386 as long as the disk controller supports the drive. Controllers supporting other drive interfaces such as RLL or ESDI work as long as the controller presents a compatible interface to the bus.

Note

This section applies only to ST506 and ESDI hard disks. SCSI disks are preformatted and use a different cylinder/head/sector translation logic, thus the considerations for ST506 and ESDI disks do not apply. Use **mkdev hd** to add hard drives to the system, regardless of whether they are SCSI or standard.

MS-DOS generally does not support non-standard disks (that is, those not defined in the ROM Fixed Disk BIOS).

Some hard disks come from the factory only partially formatted (for example, the Maxtor 1140 140 Mbyte). This problem may become evident during installation, when **badtrk**(ADM) indicates that every sector past a certain cylinder/head location is bad. You should contact the manufacturer to determine whether or not the disk is completely formatted. There are several products available that format hard disks.

Compatible hard disk controllers are discussed in the next section.

- For a standard disk

  the motherboard ROM must have an entry for the disk (type) determined by the number of heads, cylinders, tracks per cylinder (heads), sectors per track, and other characteristic information.

  Follow the manufacturer's instructions to set switches or configuration.

- For a nonstandard disk

  the user can type in information that overrides the ROM disk configuration information during installation.

  If you are unsure of what parameters to enter for your nonstandard disk, contact your disk manufacturer for this information. The **dkinit** program (called during installation) allows you to input the disk parameters.

A.6.7.1 Disks Larger than 1024 Cylinders

SCO UNIX System V/386 supports disks with more than 1024 cylinders with the following restrictions:

- The disk controller must support disks with more than 1024 cylinders.
- The disk cannot have an MS-DOS partition installed.
- If the disk is configured as the primary hard disk, the root filesystem must lie within the first 1024 cylinders. This is because the ROM BIOS cannot access boot information if it lies beyond the 1024th cylinder.

You can use the rest of the disk for swap space and/or additional filesystems. If you install only a UNIX partition, you should make certain that the root filesystem falls within the first 1024 cylinders.

For example, with the DPT PM3011 controller, SCO UNIX System V/386 works with disks containing 2048 cylinders, 16 heads, and 63 sectors, for a total maximum disk capacity of 1 gigabyte.

If the boot information lies beyond the 1024 cylinder boundary, then the following error message is generated while booting, shortly after the "loading .text" appears:

```
error in text
```

If this happens, you must rearrange your filesystems to fall within the guidelines discussed earlier.

A.6.8 Disk Controllers and Host Adapters

SCO UNIX System V/386 supports the use of ST506/ESDI hard disk controllers, and SCSI host adapters in the following combinations:

- one or two ST506 controllers (on Microchannel machines, only one is supported)
- one ESDI controller
- one or two SCSI adapters (second adapter requires manual configuration as described in "Using a Second SCSI Adapter" under "Before Installing Your Software") in the main body of these *Release Notes*
- one ST506 controller with a single SCSI adapter
- one ESDI controller with a single SCSI adapter

Note that with an ST506-SCSI mix, SCO UNIX System V/386 must be installed on the ST506. This means that if you start with SCO UNIX System V/386 installed on a SCSI disk and you want to add an ST506 disk, you must reinstall SCO UNIX System V/386 using the ST506 as the root disk.

Note that the same is true with an ESDI-SCSI mix under micro-channel: SCO UNIX System V/386 must be installed on the ESDI.

If you are running on industry standard architecture and have only SCSI disks installed on your system, you must run your computer's setup program and set the computer up for operation *without* a hard disk before installing the operating system. This forces the computer to recognize the SCSI adapter.

The following table summarizes the number of devices supported per controller/adapter:

Controller or Adapter	**Maximum Number of Devices Supported**
ST506	2 disks per controller
ESDI	2 disks per controller
SCSI	7 devices per adapter 4 disks per system 4 tape drives per adapter 4 CD-ROM drives per adapter

Note that a maximum of 4 SCSI disks are supported, regardless of the number of adapters installed. In addition, there is a limit of 7 devices per adapter. For example, if there are already 4 SCSI disks installed on an adapter, only three other devices (tape or CD-ROM) can be added.

A.6.8.1 ST506 and ESDI Controllers

Many ST506 and ESDI hard disks work with SCO UNIX System V/386. Whether or not a disk works depends upon the disk controller board used. For the disk to work, the controller must meet two tests:

1. The disk controller must be fully compatible with the standard controller for that configuration.

2. No special vendor software is needed to make the controller work under MS-DOS.

If a controller meets these tests, it should work, but if it fails these tests, it will not work.

A.6.8.2 SCSI Host Adapters

SCSI host adapters accept devices with SCSI drivers, including hard disks, tape drives, and CD-ROM drives. Each device on the SCSI bus must have an address. The address has two components: an address for the controller and an address for the device itself. The device address is known as a logical unit address (LUN). In this release, the first SCSI disk must be configured as ID-0:LUN-0 (controller address 0, device address 0). The ID number corresponds to the jumper or switch settings on the disk, so care must be taken that hard disk settings are correct. Additional disks may be any other ID between 1-7. Note that the host adapter uses one of the ID numbers (typically number 7). The LUN number must always be 0.

A.6.9 CD-ROM Drives

SCO UNIX System V/386 also supports the following CD-ROM drives configured as SCSI devices:

Toshiba XM-3201B
Sony CDU-6110-01*
NEC CDR-77

* The Sony bus DIP switch should be set to the factory configuration; see the manual supplied with the drive.

The **mkdev high-sierra** command adds support for CD-ROM filesystems to the UNIX kernel. The **mkdev cdrom** command must then be used to configure a CD-ROM drive. To bring the CD-ROM drive online, be sure and insert a disk. If you attempt to bring up the drive without inserting a disk, the message "cannot open" is displayed.

A.6.10 Mice and Other Graphic Input Devices

The following graphic input devices are supported:

Logitech Serial Mouse
Microsoft Serial Mouse
Mouse Systems PC Mouse
Mouse Systems PC Mouse II*
Microsoft Busmouse or InPort Bus Mouse
Logitech Bus Mouse
Olivetti Bus Mouse
IBM Personal System/2 Mouse
Summagraphics Bitpad

* The Mouse Systems PC Mouse II is an optical mouse. You must use it in Microsoft Serial Mouse emulation mode and add it to the system as a Microsoft Serial Mouse.

A.6.11 Using Interrupts Otherwise Allocated to COM Ports

A situation may arise where you need additional interrupts for configuring additional hardware. Interrupts 3 and 4 are normally reserved for COM2 and COM1, respectively. If you are not using these ports, or if you wish to forsake them in favor of another device, do the following:

1. If possible, you should disable the hardware presently configured at the interrupt you wish to use. You can do this by either:

   - removing the jumper used to select the interrupt (AT architecture machines)
   - using the hardware configuration disk supplied by the manufacturer to change the interrupt vector setting (MC architecture machines).

   If this is not possible, remove the card entirely.

   Although the steps that follow remove the kernel configuration, if the physical configuration is not altered, conflicts can still occur when two pieces of hardware are configured for the same interrupt.

2. Log in as *root*.

3. Edit */etc/conf/sdevice.d/sio*. This file controls whether or not the COM devices are configured into the kernel. The first entry is for COM1 and the second for COM2. Change the appropriate entry from "Y" to "N".

4. If you are reallocating both interrupts, you should completely disable the sio driver by editing */etc/conf/cf.d/mdevice* and removing the "r" in column three of the sio entry. The r indicates that the driver should be configured into the kernel (see **mdevice**(F)).

5. Disable */dev/tty1A* or */dev/tty2A* as desired. For example, the following command disables */dev/tty1A*:

   disable tty1A

6. If you actually have a multiport card installed and you did not remove the sio driver, you need to prevent the driver from identifying your card. You can do this by commenting out the entry that matches the brand of your card in */etc/conf/pack.d/sio/space.c*. Simply insert a "/*" (slash-asterisk) on a line above the entry and a "*/" (asterisk-slash) on a line below it.

7. Relink the kernel using the **sysadmsh** System→Configure→Kernel→Rebuild selection. Follow the prompts; respond y when asked if you wish to have the new kernel boot by default, and again when asked if you wish to rebuild the kernel environment.

8. To boot the new kernel, use the **sysadmsh** System→Terminate selection to shut the system down. Press ⟨Return⟩ when the reboot message is displayed.

You can then allocate the interrupts as you wish and no conflicts should occur. If you need to use either COM port again for a serial card, undo the changes described above.

A.6.12 Modems and Autodialing

Any 100% Hayes-compatible modem works using **uucp**(C) and **cu**(C). The default autodialer is for the Hayes Smartmodem 1200. We strongly recommend external modems. Autodial programs are also supplied for the Racal Vadic 3451, the 212, the Hayes Smartmodem 1200 and 2400, the Telebit Trailblazer, and the MultiTech Multimodem 224EM. Other autodialing modems can be supported by writing a dialer program, or modifying an existing one in */usr/lib/uucp*.

In addition, the new Honey DanBer UUCP package includes a large number of dialers in the *Dialers* file. However, dialer binaries are preferred for greater reliability.

See the "Building a Remote Network with UUCP" chapter of the *System Administrator's Guide* for more information on writing other dialer programs.

A.7 Standard Architecture Hardware Notes

The following sections explain what hardware can be used with 386 machines based on Standard Architecture.

A.7.1 Network Cards

SCO UNIX System V/386 supports the following boards for Ethernet on standard architecture computers:

- 3Com Etherlink 3c501 card
- 3Com Etherlink II 3c503 card
- Western Digital WD8003E card

The driver installation software for these cards allows you to install up to three boards of the same type. During installation the software device drivers are configured for various parameters for each card, such as the interrupt vector and base I/O address. These parameters must be chosen with care, so as not to conflict with other hardware in your system. Be aware of potential conflicts before setting jumpers on the cards. Note that a network board strapped at location IRQ2 should use number 9.

Install your network card using Interrupt Vector 2, the default indicated by the prompts for configuring your network card during the SCO UNIX System V/386 installation. The 503 card has a software-selectable interrupt vector; when you select the default of 2, the hardware is configured automatically. The default manufacturer's interrupt vector setting for the WD8003E and 3c501 cards is 3 (IRQ3). You must consult the documentation for your network card and jumper the board to use Interrupt Vector 2 (IRQ2) before installing SCO UNIX System V/386. If this is unchanged, this results in a conflict between the second serial port (tty2a/COM2) and the network card. If you need to use Interrupt Vector 3, you must disable the second serial port. This can be accomplished by editing the */etc/conf/sdevice.d/sio* file after completing the installation. You must change the "Y" found in the second field on the second line to "N".

When using NFS with a 3c501 card, a **mount**(NADM) of a remote filesystem must include the options **rsize=1024** and **wsize=1024** on the **mount** command line. If you use **mnt**(C), these options should be added to the filesystem entry in */etc/default/filesys*. The value of 1024 must be used due to a known problem with this card.

Note that a 3Com 3c503 (vector 3) card may appear as an ARNET 2 port serial card. This is a known problem and is of no consequence.

A.7.1.1 Graphics-Network Card Conflicts

Certain graphics adapters conflict with the addresses used by the 3c503 and WD80003E cards. In particular, the Orchid Designer and Genoa SuperVGA boards conflict with these network cards. This problem involves memory-mapped I/O.

A.7.2 Memory Cards

We have used the following memory cards:

AMI SMART PACK 2
AST
JustRAM/AT 8Mbyte Card (Monolithic Systems)
Quadram
Tecmar
Talltree Systems
Silicon Valley Systems
STB Rio Grande
Micron Chessmate

A.7.3 Serial Cards

The following serial I/O boards are supported by the SCO UNIX System V/386 for Industry Standard Architecture computers:

AMI lamb 4 and 8 port
Arnet Controls 2, 4 and 8 port (clock option not supported)
Arnet Twin port
AST FourPORT/XN
Comtrol Corporation Systems Hostess (4 port)
Comtrol Corporation Systems Hostess/550 (8 port)
CTC Versanet 4AT (4 port) and 8AT (8 port)
Digiboard 4 and 8 port
IBM standard COM1 and COM2
Kimtron Quartet 4 port*
Olivetti RS232C Multiport Board
Quadram QuadPort™
Stargate Technologies OC4400 (4 port) and OC8000 (8 port)
Tandon Quad Serial Card
UNISYS 4 port

* The Kimtron board does not work with all 386 machines.

We strongly recommend that multiport cards use 16450 serial I/O chips instead of the slower 8250 chips. If you see a "double echo" problem, it is due to slow serial I/O hardware.

Serial Card Addresses and Notes

Physical Port	Board Type	Number of ports	Primary Address Range	Alternate Address Range	Primary Status Address	Alternate Status Address
	AMI lamb**	8	0x540-0x17F	n/a	0x210	n/a
	AMI lamb**	4	0x540-0x15F	n/a	0x210	n/a
	Arnet	8	0x100-0x13F	0x280-0x2BF	0x142	0x2C2
	Arnet	4	0x100-0x11F	0x280-0x29F	0x142	0x2C2
	Arnet	2	0x100-0x10F	0x280-0x28F	0x142	0x2C2
	AST*	4	0x2A0-0x2BF	n/a	0x28F	n/a
	CTC†	8	0x160-0x19F	n/a	n/a	n/a
	CTC†	4	0x160-0x17F	n/a	n/a	n/a
COM1	Digiboard††	8	0x110-0x14F	n/a	0x151	n/a
	Digiboard††	4	0x110-0x12F	n/a	0x151	n/a
	Hostess	8	0x140-0x13F	0x500-0x2BF	0x680	n/a
	Hostess	4	0x140-0x11F	0x680-0x29F	0x700	n/a
	Kimtron	4	0x120-0x13F	n/a	0x8120	n/a
	Olivetti†††	4	0x2A0-0x2BF	n/a	0x2BF	n/a
	Quadram***	5	0x280-0x2CF	n/a	0x2D3	n/a
	Quadram***	1	0x280-0x28F	n/a	0x2D3	n/a
	Stargate	8	0x290-0x2CF	n/a	0x2D0	n/a
	Stargate	4	0x290-0x2AF	n/a	0x2D0	n/a
	Tandon***	4	0x2A0-0x2BF	n/a	n/a	n/a
	AMI lamb**	8	0x2C0-0x2FF	n/a	0x212	n/a
	AMI lamb**	4	0x2C0-0x2DF	n/a	0x212	n/a
	Arnet	8	0x180-0x1BF	0x300-0x33F	0x1C2	0x342
	Arnet	4	0x180-0x19F	0x300-0x31F	0x1C2	0x342
	Arnet	2	0x180-0x18F	0x300-0x30F	0x1C2	0x342
	AST*	4	0x1A0-0x1BF	n/a	0x1BF	n/a
	CTC†	8	0x218-0x257	n/a	n/a	n/a
	CTC†	4	0x218-0x237	n/a	n/a	n/a
COM2	Digiboard††	8	0x210-0x24F	n/a	0x250	n/a
	Digiboard††	4	0x210-0x22F	n/a	0x250	n/a
	Hostess	8	0x200-0x1BF	0x580-0x33F	0x587	0x707
	Hostess	4	0x200-0x19F	0x700-0x31F	0x587	0x707
	Kimtron	4	0x2E0-0x2FF	n/a	0x82E0	n/a
	Olivetti†††	4	0x1A0-0x1BF	n/a	0x1BF	n/a
	Quadram***	5	0x288-0x2D7	n/a	0x2DB	n/a
	Quadram***	1	0x288-0x297	n/a	0x2DB	n/a
	Stargate	8	0x190-0x1CF	n/a	0x1D0	n/a
	Stargate	4	0x190-0x1AF	n/a	0x1D0	n/a
	Tandon***	4	0x280-0x29F	n/a	n/a	n/a
OTHER1	UNISYS****	4	0x000-0x03F	n/a	n/a	n/a
OTHER2	UNISYS****	4	0x008-0x047	n/a	n/a	n/a

* Only enhanced mode is supported. Do not use the AST driver.

** Only continuous mode is supported.

*** Tandon is the only card whose I/O addresses are potentially identical with other supported cards, such as the AST and Quadram serial cards.

Because card addresses must not overlap in the same systems, if you have both a Tandon and a Quadram, the Tandon must be on COM1 and the Quadram must be on COM2.

You can use the following Quadram serial expansion cards in these configurations under SCO UNIX System V/386:

COM1	COM2
5-port	-
1-port	-
-	5-port
-	1-port
5-port	1-port

If you need to use any other configuration of Quadram cards, you must relink the kernel using the Link Kit. You also need to edit */etc/conf/pack.d/sio/space.c* as follows:

1. This file contains several groupings of serial card descriptions. Each group is identified by the first number on the description line.

   Find the descriptions for the Quadram cards in each group and move that descriptor line to the front of the group.

2. Relink and boot the new kernel. Refer to "Tuning System Performance" in the *System Administrator's Guide* for information on relinking the kernel.

**** These serial cards only work on the UNISYS PC/IT.

† Notes for the CTC Versanet serial cards:

1. The correct switch settings for the 8AT and 4AT are:

   As a COM1 (strapped at addr 0x160, using irq4) the 8AT has:

   switches 33, 35, 36, 38, 39 & 40 OFF
   switches 34, 37 ON (that is, shunted)

   on the DIPSWITCH selection:

   5, 6 & 8 should be OFF
   all the others should be ON

   As a COM2 (strapped at addr 0x218, using irq3) the 8AT has:

   switches 33, 35-40 OFF
   switch 34 ON (that is, shunted)

   on the DIPSWITCH selection:

   1, 2 & 7 should be OFF
   all the others should be ON

   The 4AT is the same as the 8AT in both the above cases, with the following common exception:

   switches 39 & 40 MUST BE ON (shunted)

2. These boards come in both 8250 and 16450 versions. You must have the 16450 version.

3. The original CTC Versanet boards used different addresses. Please ask your hardware vendor for versanet boards strapping at the above addresses (0x160 and 0x218).

4. The CTC "Maomao-4" serial board is not supported.

†† Notes on the Digiboard serial card:

1. Ports for switches DS2 to DS5 (DS9 for 8 port version) must be strapped starting at the boards base address as given in the table and incrementing by 8 for each port.

   The following example is for COM1 at 110:

DS2	110
DS3	118
DS4	120
DS5	128
DS6	130
DS7	138
DS8	140
DS9	148
DS1	150

2. If COM1 is used, then all the ports must be strapped as "EVEN" to interrupt request line 4 (see Digiboard documentation). If COM2 is used, then all the ports must be strapped as "ODD" to interrupt request line 3 (same notation).

3. Only rev. C and later boards are supported.

††† Notes for the Olivetti RS232C Multiport board:

1. The factory settings do not function properly. You must alter the existing switch positions to reflect those listed in item 4 below.

2. If you are using a single Olivetti board, you must configure it as COM2.

3. If you are using two Olivetti boards, one must be configured as COM2 and the other configured as COM1, with the COM1 port built into the M380 disabled. To disable the COM1 port built into the M380, refer to the section on "Setting Up the System" in your Olivetti *Installation and Operations Guide*.

4. The correct switch settings:

   As a COM1 (strapped at addr 0x2A0):

IRQ2	IRQ3	IRQ4	IRQ5	IRQ6	IRQ7	XA1	XA0	INT	SHR
off	off	on	off	off	off	off	on	on	off

   As a COM2 (strapped at addr 0x1A0):

IRQ2	IRQ3	IRQ4	IRQ5	IRQ6	IRQ7	XA1	XA0	INT	SHR
off	on	off	off	off	off	off	off	on	off

A.7.4 Add-On Hard Disks

Many hard disks can be used with SCO UNIX System V/386. In the case of ST506/ESDI drives, the drive must be supported by the ROM BIOS, or the ROM parameters must be entered at installation time. Check your computer hardware reference manual for the appropriate ROM table entries for your computer.

The following hard disk drives were tested and found to work correctly with SCO UNIX System V/386:

CDC Wren IV
CDC Wren V
Conner CP-340
Conner CP-3100
Quantum Q250
Quantum Q280
Quantum P40S
Quantum P80S
Syquest SQ555

The following hard disk drives are reported to work with SCO UNIX System V/386:

Maxtor XT-4380S
Maxtor LXT-100S

For disks using the Adaptec AHA-154x controller, the first disk drive (ID-0:LUN-0) should have no jumpers on A0, A1 and A2. The second disk drive (ID-1:LUN-0) should have a jumper on A0 and no jumpers on A1 and A2.

A.7.5 Compatible Hard Disk Controllers

We have used the following controllers under SCO UNIX System V/386:

Adaptec ACB-2320
Adaptec ACB-2322*
Adaptec ACB-2370 RLL
Adaptec ACB-2372 RLL*
Adaptec AHA-154x SCSI Host Adapter
Adaptec 4525 SCSI EDSI Disk Controller
SMS OMTI 8620 (ESDI controller)
SMS OMTI 8627 (ESDI controller)
Western Digital WD1010 or compatible
Western Digital WD 1003**
Western Digital WD 1003-WA2**
Western Digital WD 1005
Western Digital WD 1007
Western Digital WD 7000
DPT PM3011***
DPT MX3011***
DTC WD1010 compatible
DTC WD2010 compatible

* These models are hard disk and floppy disk controllers.

** The Western Digital controller cards use either WD1010, WD2010, or WD2020 controller chips.

*** The DPT controller cards must have EPROM 2E or later to run SCO UNIX System V/386 on a Compaq.

A.7.6 Video Adapters

This section concerns video adapters supported under SCO UNIX System V/386 for Industry Standard Architecture machines. The following table lists tested graphics adapters.

Tested Graphics Adapters

Card	Resolution	Type
Compaq Plasma	640x400	2-color
Compaq VGC	640x350	16-color
Compaq VGC	640x480	16-color
Genoa SuperVGA†‡	640x350	16-color
Genoa SuperVGA†‡	640x480	16-color
Genoa SuperVGA†‡	800x600	16-color
Genoa SuperVGA†‡	1024x768	16-color
Hercules Monochrome Graphics Card	720x350	Monochrome
IBM EGA	640x350	16-color
IBM VGA	640x350	16-color

(Continued on next page)

Tested Graphics Adapters *(Continued)*

Card	Resolution	Type
IBM VGA	640x480	16-color
Orchid Designer*‡	640x350	16-color
Orchid Designer*‡	640x480	16-color
Orchid Designer*‡	800x600	16-color
Orchid Designer*‡	1024x768	16-color
Paradise VGA Plus	640x350	16-color
Paradise VGA Plus	640x480	16-color
Paradise VGA Plus	800x600	16-color
STB Extra-EM†	640x350	16-color
STB Extra/EM†	640x480	16-color
STB Extra/EM†	800x600	16-color
STB Extra/EM†	1024x768	16-color
Video 7 Fastwrite VGA	640x350	16-color
Video 7 Fastwrite VGA	640x480	16-color
Video 7 Fastwrite VGA	800x600	16-color
Video 7 VEGA (EGA)	640x380	16-color
Video 7 VRAM VGA	640x350	16-color
Video 7 VRAM VGA	640x480	16-color
Video 7 VRAM VGA	800x600	16-color
Video 7 VGA 1024i	640x350	16-color
Video 7 VGA 1024i	640x480	16-color
Video 7 VGA 1024i	800x600	16-color

* The Orchid graphics adapter card does not work in extended mode on a 20 megahertz or faster bus. In addition, the Orchid uses IRQ2 during certain refresh operations in particular video modes.

† If you are using STB VGA/EM or Genoa Super VGA graphics boards, character mode is not properly restored to the console when exiting. Certain characters are followed by vertical bars when you terminate SCO UNIX System V/386.

‡ The Orchid Designer and Genoa SuperVGA boards conflict with the addresses used by the 3c503 and WD80003E network cards. This problem involves memory-mapped I/O.

The following table includes graphics adapters that are untested but reported to work:

Untested Graphics Adapters

Card	Resolution	Type
EIZO MD-B07	640x350	16-color
EIZO MD-B07	640x480	16-color
EIZO MD-B07	800x600	16-color
EIZO MD-B10	640x350	16-color
EIZO Extra/EM	640x480	16-color
EIZO Extra/EM	800x600	16-color
EIZO Extra/EM	1024x768	16-color
Quadram QuadVGA	640x350	16-color
Quadram QuadVGA	640x480	16-color
Quadram QuadVGA	800x600	16-color
Quadram QuadVGA	1024x768	16-color
Tecmar VGA/AD	640x350	16-color
Tecmar VGA/AD	640x480	16-color
Tecmar VGA/AD	800x600	16-color
Tecmar VGA/AD	1024x768	16-color
Toshiba Grid 758 Display	640x400	Monochrome
Toshiba Grid Plasma Display	640x400	Monochrome

Note that graphics adapter cards that are not listed in the previous tables but are identical to cards listed above should work as well.

A.7.7 Tape Drive/Controller Combinations

The tape driver included in this release works with the following drive/controller combinations:

Supported QIC-02 Tape Controllers and Cartridge Drives

Manufacturer	Controller	Drive	Type	Notes
Archive	SC400	Scorpion 5945	A	
Archive	SC402	Viper QIC-02/60MB	A	(4,5)
Archive	SC402	Viper QIC-02/150MB	A	(4,5)
Archive	SC499r	Scorpion 5945	A	
Bell Technologies	PC-36	XTC-60	W	
Cipher	QIC-02	CP-60B	W	
Cipher	QIC-02	CP-125B	W	
Cipher	Cipher 811/817	5400/plus	E	
Computone	SC400	Scorpion 5945/60MB	A	
COREtape	PC-36	5000(E)	W	
Emerald	xnx-50-2012	Cassette	E	(3)
Emerald	xnx-60-2002	Cartridge	E	(3)
Everex	PC-36	60MB Internal	X	
ITT	PC-36	5000(E)	W	
Mountain	QIC-02	60MB Filesafe	M	
Mountain	QIC-02	150MB Filesafe	M	(5)
Mountain	QIC-02	300MB Filesafe	M	(5,6)
Mountain	PC-36	60MB Internal	W	
Mountain	PC-36	150MB Internal		(5)
Olivetti	PC-36	5000(E)	W	(1)
Tecmar	PC-36	QIC-60AT	T	
Tecmar	PC-36	QT-60I	T	
Tecmar	PC-36	QT-90I	T	
Tecmar	PC-36	QT-125I	T	
Tecmar	PC-36	QT-150I	T	
Tecmar	PC Bus Host Adapter	QT-60E	T	
Tecmar	PC Bus Host Adapter	QT-90E	T	

(Continued on next page)

Supported QIC-02 Tape Controllers and Cartridge Drives
(Continued)

Manufacturer	Controller	Drive	Type	Notes
Tecmar	PC Bus Host Adapter	QT-122E	T	
Tecmar	PC Bus Host Adapter	QT-150E	T	
TI	SC400	Scorpion 5945/60MB	A	
Wangtek	PC-36	5000(E)	W	
Wangtek	PC-36	5xxxPK/125,150MB	W	
Wangtek	PC-36	5xxxEK/60,125,150MB	W	

Most controllers that conform to the QIC-02 standard should work, but only the units listed above have been tested.

In addition, the following floppy tape drive/controller combinations are supported:

Supported Floppy Tape Controllers and Cartridge Drives

Manufacturer	Drive	Notes
Alloy	APT-40/Q (40 Mbyte)	
Irwin	110 (10 Mbyte)	(2,3)
Irwin	125 (20 Mbyte)	(2,3)
Irwin	145 (40 Mbyte)	(2,3)
Irwin	185 (80 Mbyte)	(2,3)
Irwin	225 (20 Mbyte)	(2,3)
Irwin	245 (40 Mbyte)	(2,3)
Irwin	285 (80 Mbyte)	(2,3)
Irwin	2020 (20 Mbyte)	(2,3)
Irwin	2040 (40 Mbyte)	(2,3)
Irwin	2080 (80 Mbyte)	(2,3)
Mountain	TD44-40 (40 Mbyte)	
Wangtek	FAD 3500 (40 Mbyte)	
Tecmar	QT-40i (40 Mbyte)	

Default Settings

Manufacturer	Base Address	DMA Channel	Interrupt	Type
Archive	0x220	3	3	A
Compaq	0x300	1	5	W
Computone	0x200	1	4	A
Emerald	0x300	3	2*	E
Everex	0x2C0	1	5	X
ITT	0x338	1	5	W
Mountain	0x28C	1	3	M
Tecmar	0x330	1	5	T
TI	0x220	3	3	A
Wangtek	0x338	1	5	W

* When installing a tape drive using interrupt vector 2, set the hardware for IRQ2 and supply interrupt vector 25 when running **mkdev tape**. The hardware bootup message will indicate an interrupt vector of 31.

Notes:

1. The Olivetti tape controller's factory jumper settings do not match the type W drive default values. To use the Olivetti drive without changing the controller card's jumper settings, use the **mkdev tape** command to select all four of the following settings:

Type	DMA	Interrupt	Base Address (hex)
W	1	25	0x288

2. The Irwin 40 Mbyte drive can use either DC2000 or DC1000 tape cartridges and the 10 and 20 Mbyte drives use a DC1000 tape cartridge. A DC1000 tape written on a 125 or 145 model drive can be read by either drive.

3. Irwin and Emerald drives now function properly on 20-MHz 386 machines.

4. Archive drives using the SC402 controller do not use the default type A interrupt 4. Use the **mkdev tape** command to change the interrupt or restrap the controller card.

5. High density 150 Mbyte drives require DC6150 or DC600XTD tapes for writing. Low density tapes can be read but not written.

A.7.8 SCSI Tape Drives

This section discusses SCSI tape drives.

Certain SCSI tape drives are easier to use if a "no unload" device is used:

- Some DATs and Exabyte drives take as long as a minute to load. Use of the "no unload" device prevents the tape from being automatically unloaded (ejected) after each access.

- Certain tape drives (such as the Wangtek 5525ES) also require a load command after an unload. The actual unload is not obvious because the tape is not physically ejected as it is with DAT drives. (The unload is done on a **close()** after a read or write sequence.)

The following SCSI tape drives are supported:

Supported SCSI Tapes Drives

Manufacturer	Drive	Notes
Archive	60/150 Mbyte	(1)
Archive Python DDS 4520NT	DAT	
Archive Python DDS 4521NT	DAT	
Exabyte	8mm Cartridge	
Hewlett-Packard 35450A	DAT	(2)
Mountain	8mm Cartridge	
Wangtec 6130FS (full height)	DAT	(2)
Wangtec 6130HS (half height)	DAT	(2)
Wangtek 5525ES	1/4 inch	

Notes

1. External unit has two tape slots, each with 150 Mbyte capacity, giving this tape drive unit a 300 Mbyte capacity.

2. Wangtek and Hewlett-Packard DAT drives require a **tape erase** command before the first write to a new tape.

A.7.8.1 DAT Drives

The following notes apply to DAT drives.

- The first command to a tape after power-up of the system will fail. There will be no error message. After a tape change, some devices issue a warning on the next tape access in the form of "Unit attention" or a longer message for SCSI-II devices. In either case, a tape command will have to be repeated to ensure success. However, a read or write from/to the tape will succeed if there are no other errors.
- You should wait for the drive to finish its initialization sequence before attempting to access the device. This is usually very swift on a non-DAT device but may take 30 seconds or more for a DAT.
- If you make more than one read/write access to a tape, it is quicker to use the no unload device and just press the unload button on the drive to retrieve the tape.

A.7.9 Typical Device Interrupts

Interrupt	Device
0*	Clock
1*	Console
2	Networks, Tapes and others
3	Serial COM2
4	Serial COM1
5	Alternate Parallel Port (lp2)
6*	Floppy Disk
7	Main Parallel Port (lp0 or lp1)

* Do not use these interrupts.

A.7.10 SCSI Guidelines

SCO UNIX System V/386 supports up to two AHA-154x or WD7000 host adapters. Be sure and consult the "Hard Disk Controllers and Host Adapters" section under "General Compatibility Guidelines" for additional information about SCSI host adapters and devices.

A.7.10.1 Western Digital WD7000 Notes

The WD7000 adapter has not been tested extensively and certain hardware configurations do not appear to work. If the factory configuration is used, the system typically hangs during the installation of the N2 diskette. Machines that have been tested or reported to work with the WD7000-ASC are listed later in this section. (Note that other systems not listed here may also work.) If you are using a machine that is not on this list, or is on the list of machines that have problems, the system usually reboots unexpectedly during normal operation.

The following describes the correct configuration for the WD7000-ASC:

Interrupt	12 decimal (octal 14)*
DMA	6
Base address	0x350-0x352

* The switches are in decimal and the bootup messages are in octal.

The jumper settings are shown in the four tables that follow (only the significant banks are shown).

Note

If you are using the WD7000 as the controller for floppy disks, remove the jumpers on W6 and W9. If you have a separate floppy controller, place a jumper on W6.

W2 (2,5,9 jumpered)

:	[:]	:	:	[:]	:	:	:	[:]	:	:
1	2	3	4	5	6	7	8	9	10	11

W3 (1,3,5 jumpered)

[:]	:	[:]	:	[:]
1	2	3	4	5

W1 (none jumpered)

:	:	:	:	:
1	2	3	4	5

W4 (4 jumpered)

:	:	:	[:]	:
1	2	3	4	5

If these jumper settings do not solve the problem, check the following:

1. The first and last devices on the bus must be terminated:

    - If the adapter is the first device, there should be resistors just behind the connector
    - If the only other device is the hard disk, this is therefore the last device and therefore must be terminated
    - We recommend that if there are problems all other devices should be removed from the SCSI bus.

2. The first physical SCSI hard disk must be configured as ID 0.

3. Parity checking should be on.

The following systems have been reported to work with the WD7000, but we have not tested them:

ACER 386/20
ALR Flexcache 25386 (w/IOCR mod on 7000)
AMI Mark II
ARC Skyscraper
AST Premium 386/16
AST Premium 386/25 (w/IOCR mod on 7000)
AT&T 6386WGS
Compaq Deskpro 386/25
CSS Mother Superior

Everex Step 25
Intel
Micronics
Mitac 25MHz
NCR PC-920
NCR PC-925
Rose Hill 387
Tandy 3000
Tandon 386/20
Unisys 386/25
Unisys PC-6
TI 386/20
Wyse 16MHz
Wyse 3216-40

Difficulties or incompatibilities have been reported with the following machines:

Compaq 386/20E
PC Kraft 386/25 with 387/20
Unisys 386/20

A.7.10.2 Adaptec AHA-154X Notes

Not all manufacturer's computer systems are compatible with Adaptec host adapters. If you are unable to get the host adapter to function or experience file corruption, use the test documented in this section to see if an incompatibility exists. You can also contact your manufacturer or Adaptec to determine if your hardware is incompatible.

The Adaptec AHA-154X host adapter supports first-party DMA, asynchronous and synchronous peripherals, jumper selectable DMA and interrupt channels and I/O port addresses, programmable mailbox architecture that allows multithreaded operation on the SCSI bus, and even or odd starting address transfers. The AHA-154X back-end software allows setting of 154X adapter-configurable parameters. Some of the possible settings are:

- bus on-time
- bus off-time
- transfer rate
- number of CCBs (Command Control Blocks)
- number of mailboxes

The Adaptec SCSI host adapter should run correctly with the manufacturer's default settings. It should be configured at base address 0x330, IRQ 11, and DMA channel 5. However, on some computers, such as the Tandy 4000, you must remove jumper J5 if the host adapter is an AHA-1540 or 1540a, or jumper J8 if the host adapter is an AHA-1542. (For more information, refer to your Adaptec documentation.)

If you are installing a second Adaptec AHA-1540 board, you must reset several jumpers to be configured at base address 0x230, IRQ 12, and DMA channel 6. Some computers, such as the Tandy 4000, require you to remove jumper J8. For two AHA adapters to coexist, the BIOS must be disabled on the second adapter. This can be done with the newer AHA boards (1540a/1542a) by removing a jumper, but not with the older 1540 boards. You cannot use two 1540 boards in one machine. (For more information, refer to your Adaptec documentation.)

A.7.10.3 Formatting and Verifying Devices

The Adaptec 154x series of controllers have a set of BIOS routines to maintain and debug SCSI devices.

First you will need to know what BIOS address is on the Adaptec controller. This information can be found in the Adaptec manual. The default address is 0DC000H. Once you have this number you can perform a number of tasks. One such task is the formatting and verifying of SCSI devices. All SCSI devices (hard disks) should be formatted prior to installing SCO UNIX System V/386.

To format your devices you will need to do the following:

1. Boot DOS
2. Run the DEBUG program
3. At the "-" prompt enter: **g=dc00:6** (notice how the BIOS address is truncated). A menu is displayed.
4. Select the format option on the drive specified. (Note that this destroys all data on that disk.)
5. When the format is complete, select the verify option on the drive specified.
6. Exit the menu.

A.7.10.4 Hardware Incompatibilities

There is an option that will allow you to test the interface between the Adaptec controller and motherboard. To run this test, do the following:

1. Boot DOS
2. Run the DEBUG program
3. At the "-" prompt enter: **g=dc00:9** (notice how the BIOS address is truncated).

You should see a set of diagnostic messages displayed as the test is run. If you see a fail message, or no diagnostic messages are displayed, a hardware incompatibility exists. Your system may have trouble loading and running SCO UNIX System V/386. Since this is a hardware level inconsistency, there is no software patch available.

A.8 Microchannel Architecture Hardware Notes

The following sections explain what hardware can be used with 386 machines based on Microchannel Architecture.

A.8.1 Network Cards

SCO UNIX System V/386 supports these Ethernet boards:

- 3Com Etherlink/MC 3c523 card
- Western Digital WD8003E card

The driver installation software for these cards allows you to install up to three boards of the same type. During installation the software device drivers are configured for various parameters for each card, such as the interrupt vector and base I/O address. These parameters must be chosen with care, so as not to conflict with other hardware in your system. Be aware of potential conflicts before changing the parameter values of the cards using the hardware configuration disk supplied by the manufacturer. Note that a network board strapped at location IRQ2 should use number 9.

You should install your network card using Interrupt Vector 2, the default indicated by the prompts for configuring your network card during the SCO UNIX System V/386 installation. The default manufacturer's interrupt vector setting for the WD8003E and 3c523 cards is 3 (IRQ3). You must consult the documentation for your network card and reconfigure the board to use Interrupt Vector 2 (IRQ2) using the hardware configuration disk supplied by the manufacturer, before installing SCO UNIX System V/386. If this is unchanged, this results in a conflict between the second serial port (tty2a/COM2) and the network card. If you need to use Interrupt Vector 3, you must disable the second serial port.

This can be accomplished by editing the */etc/conf/sdevice.d/sio* file after completing the installation. You must change the "Y" found in the second field on the second line to "N".

A.8.2 Serial Cards

The following serial cards are supported under SCO UNIX System V/386 for Microchannel Architecture machines:

IBM PS/2 model 3033 dual async 2 port card (up to 3 per system)
Stargate PLUS 8 MC
AST 4 port or 8 port Async Cluster Adaptor
Digiboard PS-COM/8 port or 16 port
Arnet Multiport/2 port or 8 port - 1 or 2 cards
Comtrol Hostess/MC 4 port
Comtrol Hostess 550/MC 8 port

Serial Card Addresses and Notes

Physical Port	Board Type	Base Number of ports	Address
COM1	Motherboard	1	0x3F8
COM2	Arnet*	16	0x140
	IBM	2	0x2F8
	AST	4	0x2F8
	Stargate**	8	0x400
	Digiboard	16	0x3000
	Digiboard	8	0xDB80
	HostessMC	4	0x500
	HostessMC	8	0x500
	HostessMC	4	0x540
	HostessMC	8	0x540
	HostessMC	4	0x580
	HostessMC	8	0x580

* You must use the reference disk to modify the Optional I/O Address Block to be 0108-018F hexadecimal.

** On the Apricot Qi, the Stargate serial adapter card clashes with the Ethernet Controller start address. To resolve the problem, use the reference disk to modify the Ethernet Controller start address. Choose the menu options in the following sequence:

Configuration
Change
Internal Ethernet Controller
Port Address
Alternate #7

Dumb boards on COM 2 such as the Hostess, AST 4 port, and the IBM Dual Async occupy addresses as shown in the table below:

Port Number	Address
1	0x2F8
2	0x3220
3	0x3228
4	0x4220
5	0x4228
6	0x5220
7	0x5228
8	0x2F0

A.8.3 Video Adapters and Monitors

Microchannel architecture machines have the VGA built into the mother board, therefore no options are available. Adapters that come as the default adapter in supported computers (monochrome or color monitor) work with SCO UNIX System V/386.

A.8.4 Video Cards and Monitors

The following video cards are supported under SCO UNIX System V/386:

IBM Personal System/2 Integral VGA adapter
and the 8503, 8512, 8513, and 8514 monitors
Olivetti P800, Integral VGA
Olivetti P500, Integral VGA
Tandy 5000MC, Integral VGA
Apricot Qi, Integral VGA

A.8.5 Compatible Hard Disk Controllers

The following controllers are supported for the Microchannel Architecture:

Adaptec AHA-1640 (SCSI host adapter)*
Adaptec 2610 (ESDI disk controller)
Adaptec 2620 (ST506 disk controller)
IBM (ESDI disk controller)
IBM (ST506 disk controller)
IBM (SCSI disk controller)**
Western Digital 1006V-MCI (ESDI disk controller)
Western Digital 1007V-MCI (ST506 disk controller)

* Most PS/2 clones

** IBM PS/2 only (not compatibles)

A.8.6 Tape Drive/Controller Combinations

Personal System/2 tape support is being added as Personal System/2 tape devices are becoming available. Call your software support center for the 'latest tape driver availability.

Supported Tape Controllers and Cartridge Drives

Manufacturer	Controller	Drive	Type
Everex/Archive	QIC-02	2150L	E
Archive	SCSI	2150S	A
IBM	QIC-02	6157	X
IBM	QIC-02	6157-002	X
Irwin	floppy	245	N/A
Irwin	floppy	285	N/A
Mountain 7060	QIC-02	60MB Filesafe	M
Mountain 7120	QIC-02	150MB Filesafe	M
Tecmar/Wangtek	QIC-02	QT-60E	T
Tecmar/Wangtek	QIC-02	QT-150E	T

Default Settings

Manufacturer	Base Address	DMA	Interrupt	Type
IBM 6157	0x3120	2	6	X
Mountain	0x200	1	5	M
Everex/Archive	0x300	1	3	E
Tecmar/Wangtek	0x300	3	5	T

A.8.7 SCSI Guidelines

Consult the "Hard Disk Controllers and Host Adapters" section under "General Compatibility Guidelines" for information about SCSI host adapters and devices.

Appendix B
SCO UNIX System V/386
Release 3.2 Version 2.0
Using the Korn Shell (ksh)

B.1 Introduction

The Korn Shell, **ksh**(C), is an interactive command-language interpreter and programming language that reads and executes commands from either the terminal or a file. The **ksh** combines the best features of the two common UNIX System shells, the standard Bourne shell, **sh**(C), and the C shell, **csh**(C). The **ksh** provides both compatibility with **sh** and the command history and substitution features of **csh**. In addition, **ksh** includes command-line editing, job control, and enhanced command-history functionality.

B.2 Starting ksh

If you are currently running **sh** or **csh**, you can start **ksh** by entering:

ksh

at the command line. When you run **ksh** as a program from your original login shell, you cannot automatically access the command-line editing and job control features. However, by typing:

stty susp '^Z'
set -o vi

you enable both job control and command line editing.

To use **ksh** as your default login shell, ask the system administrator to change your login-shell specifier in the */etc/passwd* file to **ksh**. When the system administrator specifies **ksh** as the login shell when creating a new user, the **sysadmsh**(ADM) utility creates two files in the user's home directory: *.profile* and *.kshrc*.

When you log in using **ksh** as your login shell, the shell reads commands from the system profile file, */etc/profile*, and then from *.profile* in the current directory or $HOME/*.profile*, if either file exists. Next, the shell reads commands from the **ksh** environment file, $HOME/*.kshrc*, if it exists. If there is no *.sh_history* file (the history file where **ksh** stores commands that you enter at the keyboard) in the your home directory, **ksh** creates one.

For more information on these files and how to modify your **ksh** environment, see the section later in this appendix, "Customizing the ksh Environment."

B.3 Using the ksh Built-in Editors

Using **csh** or **sh**, the only way to fix errors on the command line is to backspace or retype the entire line. With **ksh**, you can edit the command line using the familiar commands that you use to edit files. The **ksh** provides both **vi**-like and **emacs**-like built-in editor interfaces for editing the command line.

At login time, **ksh** reads the *.kshrc* environment file and turns on the **vi**-like editor. You can turn off the **ksh** editor functionality completely or turn off **vi** and turn on **emacs** for the current session or for each login session.

To turn off **vi** for the current login session only, enter the following at the command line:

set +o vi

To turn on **emacs** for the current login session, enter the following at the command line:

set -o emacs

To turn on or off either the **vi** or **emacs** editors automatically when you log in, add the appropriate command to the *.profile* or environment file (*.kshrc* by default).

You can also use the EDITOR and VISUAL environment variables to set the editor to any pathname that ends in **vi**. For example, to turn on the **vi** editor automatically when you log in, add the following line to your *.kshrc* file:

EDITOR=/usr/bin/vi

Note

The VISUAL variable overrides the EDITOR variable.

This appendix includes information on the built-in **vi**-like editor. For information about using the **emacs**-like editor, see your **emacs** documentation.

B.3.1 Using the vi Built-In Editor Modes

Like the **vi** text editor, **ksh**'s built-in **vi** editor has two modes: *input mode* and *control mode*. In input mode, **ksh** inserts the characters that you type at the keyboard in an editing buffer. In control mode, **ksh** interprets the characters that you enter at the keyboard as editing commands.

When you log in using **ksh** as your login shell, you are in input mode automatically. (This differs from the **vi** text editor; you are initially in control mode and you must press **a** or **i** to begin entering text.) To enter control mode from input mode, press ⟨ESC⟩. If you press ⟨ESC⟩ while in control mode, the terminal beeps.

B.3.2 Editing in Input Mode

While entering commands in input mode, you can edit the command line using editing commands from the following table:

Input Mode Editing Commands

Command	Description
⟨CTL⟩**h** or ⟨BKSP⟩	moves back one character
⟨Return⟩ or ⟨CTL⟩**m**	executes the current line
⟨CTL⟩**v**	escapes the character that follows (for entering control characters)

B.3.3 Editing in Control Mode

At any time before you press ⟨Return⟩ to execute the command, you can press ⟨ESC⟩ to enter control mode. In control mode, you can move around the command line as if you were in **vi**, editing a file.

The following table shows the **vi** commands for moving the cursor on the command line in control mode:

Moving the Cursor

Key	Description
h	moves left one character
l	moves right one character
b	moves left one word
B	moves left one word, skipping punctuation
w	moves right one word
W	moves right one word, skipping punctuation
e	moves to the last character of the next word
E	moves to the last character of the next word, skipping punctuation
0	moves to the beginning of the current line
$	moves to the end of the current line
^	moves to the first character on the current line that is not a ⟨Space⟩ or ⟨TAB⟩
f*x*	moves right to the next occurrence of *x*
F*x*	moves left to the preceding occurrence of *x*
t*x*	moves right to the character before the next occurrence of *x*
T*x*	moves left to the character following the preceding occurrence of *x*
;	repeats the last character search **f**, **F**, **t**, or **T**.
,	reverses the last character search **f**, **F**, **t**, or **T**.

The following table gives the commands for entering input mode from control mode, and for changing and deleting text:

Adding, Changing, and Deleting Text

Key	Description
a	enters input mode after the character under the cursor
A	enters input mode after the last character on the line
i	enters input mode before the character under the cursor
I	enters input mode before the first character on the line
_	appends the last word of the previous **ksh** command to the current line and then enters input mode.
r*z*	replaces the character under the cursor with *z*
R*text*	replaces characters with *text* beginning at the cursor
c*motion*	changes the characters from cursor position, using the **vi** *motion* command For example:
cw	changes word below cursor
cl	changes character below cursor and then adds text
c$	changes from the current character to the end of line
cc	deletes the entire line and returns to input mode (same as c$)
x	deletes the character under the cursor
X	deletes the character to the left of the cursor
dw	deletes the word under cursor
d*motion*	deletes characters, starting at the cursor, up to and including the other end of *motion*
D	deletes from the cursor to the end of line

Adding, Changing, and Deleting Text *(continued)*

Key	Description
d$	same as **D**
dd	deletes the entire line
y*motion*	yanks the current character using **vi** motion command
Y	yanks from cursor to end of line
y$	same as **Y**
yy	yanks the entire line into the buffer
p	puts previously yanked (or deleted) words to the right of the cursor
P	puts previously yanked (or deleted) words to the left of the cursor

The following table shows the control mode commands for executing and redrawing the current line, repeating commands, and undoing modifications on the command line:

Miscellaneous Control Mode Commands

Command	Description
⟨Return⟩ or ⟨CTL⟩**m**	executes the current line
⟨CTL⟩**l**	redraws the current line
~	changes the case of the character under the cursor
.	repeats the most recent **vi** command
u	undoes the previous **vi** command
U	undoes all modifications on the current line

B.4 Accessing Commands in the History File

Using the **vi** built-in editor, you can access previously entered commands that are stored in your **ksh** history file (*.sh_history* by default). Once you retrieve a command, you can modify it and execute it again.

B.4.1 Displaying Commands in the History File

To display the list of the commands that are stored in the history file, enter **history**. The **history** command is a predefined alias that uses the **ksh** built-in command, **fc** (fix command), to access the history file. For more information about **fc**, see **ksh**(C).

The **history** alias displays the last 16 (or fewer, if there are fewer than 16 commands in the file) commands in the history file. You can specify how many and which commands that you want **history** to display. Note that the commands must be accessible in the history file for **history** to display them. The following list gives examples of how to use the **history** alias:

history -4 displays the previous four commands only.

history 20 displays all commands from the history file, starting with 20.

history 12 24 displays only commands 12 through 24.

B.4.2 Reexecuting Previous Commands

The **ksh** also includes **r**, another predefined alias that uses the **fc** built-in command. The **r** alias allows you to reexecute commands from the history file. This alias functions similarly to the **!** command in **csh**. The following list shows some common uses of the **r** alias:

Alias	Description
r	reexecutes the last command entered
r *command*	reexecutes the last *command* entered
r *x*	reexecutes the last command beginning with *x*
r #	reexecutes command number #

Note that **r** simply reexecutes commands from the history list; **r** does not allow you to modify commands before you execute them.

B.4.3 Editing Previous Commands

You can use **vi** commands to search for and retrieve commands from the history file. Once you locate a command, you can edit and reexecute it using the **vi** commands described in the section "Editing in Control Mode" earlier in this appendix.

To move through the history file, first press ⟨ESC⟩ to enter control mode. Then, use the **vi** commands in the following table to move up and down in the history file:

Moving in the History File

Command	Description
k	moves up (previous) one command in the history file
j	moves down (next) one command in the history file
/string	searches left and up (back) through the history file for the next command containing *string*
?*string*	searches right and down (forward) through the history file for the next command containing *string*
G	goes back to the oldest accessible command in the history file
n	repeats the last / or ? search command
N	repeats the last / or ? command, searching backward

B.5 Using Job Control

The job control feature of **ksh** allows you manipulate jobs running in the foreground and background. With job control, you can stop and restart programs, and move them between the foreground and the background.

Like **sh** and **csh**, the **ksh** runs commands in the foreground by default; the shell waits for the current command to finish executing before displaying the prompt. You run a command in the background by adding an ampersand (&) character to the end of the command before pressing ⟨Return⟩. Before displaying the prompt, the shell displays the status of any completed background jobs.

When job control is active and you run a command in the background, **ksh** displays both the number of the job in square brackets, [ ], and the process ID number (PID).

Job control is active by default. If job control is disabled, **ksh** displays the following error message when you try to manipulate foreground and background jobs:

```
no job control
```

To activate job control, add the following line to your environment file (*.kshrc*):

set -o monitor

To use job control, Suspend must be set (usually ⟨CTL⟩z). To set it, add the following line to your $HOME/*.profile* file:

stty susp '^Z'

B.5.1 Referring to Jobs

When you use **ksh** as your login shell, you can refer to jobs in several different ways:

PID	refers to the job by the process ID number.
%number	refers to the job by the job number that **ksh** displays in square brackets.
%string	refers to the job whose command begins with *string*.
%?string	refers to the job whose command contains *string*.
%+ or %%	refers to the current job.
%-	refers to the previous job.

For example, if you enter the command **sleep 30&**, and **ksh** displays:

```
[1]    3456
```

you can refer to this background job in various ways, including the following:

3456
%1
%sle
%?ee

B.5.2 Using the ksh Job Control Commands

The **ksh** includes the following built-in commands for process control. These commands take a PID, job name, or number as the argument.

Job Control Commands

Command	Description
bg	starts the specified stopped job in the background.
fg	moves the specified background job to the foreground.
jobs	displays status information about current jobs.
kill	terminates the specified background job.
wait	tells **ksh** to wait for all or a specific background process to complete.

B.5.3 Running Jobs in the Background

To run a job in the background, enter the command name, followed by an ampersand (**&**) character and press ⟨Return⟩.

For example, when you enter **sleep 30&**, **ksh** starts the job in the background and displays a message like the following:

```
[1]    3456
```

When the job finishes, **ksh** displays a message like this:

```
[1] + Done                    sleep 30&
```

B.5.4 Moving Background Jobs to Foreground

To move a job that is running in the background to the foreground, enter **fg** followed by the PID, command name, or job number.

For example, move the **sleep** background process to the foreground by entering:

fg %?ee

Once the job is in the foreground, **ksh** waits for it to complete before displaying a prompt.

B.5.5 Moving Foreground Jobs to Background

To move a foreground job to the background, press Suspend. The **ksh** stops the job and displays a message like the following:

```
[1] + Stopped                 sleep 30&
```

The **ksh** displays your prompt. At the prompt, enter **bg** followed by the PID, job number, or name. For example, move the **sleep** process to the background by entering:

bg %?ee

The **ksh** restarts the suspended job, displays a message like the following, and runs the job in the background:

```
[1]   sleep 30&
```

B.5.6 Displaying Information about Jobs

Use the **jobs** command to display status information about the jobs that **ksh** is currently running. To display information about all active jobs, enter **jobs** at the command line. The **ksh** displays status information in the following format:

```
[2] + Running                 sleep 40&
[1] - Stopped                 sleep 30&
```

The **ksh** displays a plus (+) character after the job number of the current job, a minus (-) after the previous job.

Use the **-l** option to display the PID after the job number. For example, if you enter **jobs -l**, **ksh** displays information like the following:

```
[1] + 23909  Stopped            sleep 30&
```

You can limit the display to PID only by using the **-p** option to **jobs**.

B.5.7 Terminating a Background Job

To terminate a process, use the **kill**(C) command. The syntax for **kill** is:

kill [*-signal*] *job*

where *signal* is an optional signal number or name and *job* is the job name, number, or PID of the job that you want to terminate.

To display a complete list of signal numbers and names, enter **kill -l**. If you do not specify *signal*, **ksh** sends the TERM (terminate) signal to the specified job. See the previous section "Referring to Jobs" in this appendix for more information on referring to jobs by job name, number, and PID.

When you use **kill**, **ksh** displays a message like the following and terminates the job:

```
[1] + Terminated        sleep 30&
```

B.6 Customizing the ksh Environment

The **ksh** uses two files located in your home directory, *.profile* and *.kshrc*, to set up your environment. Use the *.profile* file to set variables and options for your login shell and other programs that you run from the **ksh**. The *.kshrc* file is the **ksh**-specific environment file; use it to define aliases and set **ksh** command-line options.

B.6.1 Modifying the .profile File

Whenever you log in using **ksh** as your login shell, **ksh** executes the *.profile* file, automatically executing commands and setting exported environment variables. Commands and environment variables stored in your *.profile* file must appear in the same format that they would if you were entering them at the keybord.

To assign values to **ksh** environment variables, use this format:

```
EDITOR=/bin/vi
export EDITOR
```

The following table shows some of the more common environment variables:

Environment Variables

Variable	Description
COLUMNS	specifies the number of columns that **ksh** uses to display the command line
EDITOR	sets either the **vi**-like or **emacs**-like editor to use when editing the command line
ENV	sets the environment file (*.kshrc* by default)
HISTFILE	sets an alternate history file (*.sh_history* by default)
HISTSIZE	sets the maximum number of commands that are stored in the history file (128 by default)
HOME	specifies the default argument used by the **cd**(C) command
MAILCHECK	specifies the interval in seconds that **ksh** checks for new mail (600 seconds by default)
PATH	specifies the pathnames that **ksh** searches when executing commands
PS1	specifies the primary prompt to display when the **interactive** option is on ($ by default); **ksh** replaces an exclamation point (!) with the command number (to print a ! in the prompt, enter !!)
TERM	specifies the type of terminal that you are using
VISUAL	sets the editor to use when editing command lines (overrides the value of EDITOR)

For example, to set up your prompt to include the machine name and command number, add the following lines to your *.profile*:

```
PS1='! fscott'
export PS1
```

For a complete list of environment variables used by **ksh**, see **ksh**(C).

B.6.1.1 Executing a File on Logout

You can use the **trap 0** command in your *.profile* file to instruct **ksh** to execute a file, for example, *.logout*, when you exit the shell. To do this, enter the following in *.profile*:

```
trap $HOME/.logout 0
```

The *.logout* file must be executable.

B.6.2 Modifying the .kshrc File

The *kshrc* file contains definitions for aliases and functions and default option settings for **ksh**. You should specify any commands and definitions that only **ksh** recognizes in this file rather than in *.profile*.

B.6.2.1 Defining Aliases

An alias is an abbreviated name for a command. Use the following format to define an alias:

alias *shortname='commandname'*

where *shortname* is the abbreviated name for *commandname*. For example, to define **ls** as an alias for **lc -F**, use the following command:

```
alias ls='lc -F'
```

You can define aliases for the current **ksh** session by entering **alias** at the command line. To specify that an alias definition remain in effect across login sessions, add the alias command to your *.kshrc* (or other **ksh** environment file).

You can display the complete list of your aliases by entering **alias** at the prompt. To display a particular alias definition, enter **alias** followed by the alias name. For example, if you enter **alias ls**, **ksh** displays:

```
ls=lc -F
```

To unset a particular alias (in either the *.kshrc* file or at the command line), enter **unalias** followed by the alias name.

When you run scripts that do not invoke another **ksh**, regular aliases do not remain defined. However, you can *export* these alias definitions by defining them in your environment file using this format:

alias -x ls='lc -F'

The **ksh** automatically predefines several aliases. These aliases are compiled into the shell, but you can unset or redefine them. (We do not recommend redefining preset aliases.) The two most common preset aliases are **history** (for displaying the contents of the *.sh_history* file) and **r** (for reexecuting previously entered commands). For more information about preset aliases, see the section "Accessing Commands in the History File" in this appendix and **ksh**(C).

B.6.2.2 Setting ksh Options

You can set **ksh** command-line options in your environment file. Use the following format:

set -o *option*

To unset an option, use a plus (+) character in place of minus (-).

The following table lists some useful **ksh** options:

set Options

Option	Description
allexport	exports all subsequent variables automatically (same as **-a**)
bgnice	runs all background jobs at a lower priority (set by default)
emacs	uses the **emacs**-like built-in editor for command-line editing
ignoreeof	prevents **ksh** from exiting on end-of-file, ⟨CTL⟩**d**; (when **ignoreeof** is set, you must enter **exit** to terminate the shell)
monitor	runs background jobs as a separate process group and prints status when a process completes (systems with job control set the **monitor** option automatically for interactive shells)
verbose	prints shell input lines as they are read
vi	uses the **vi**-like built-in editor for command-line editing

For a complete list of **set** options and descriptions, see **ksh**(C).

B.6.3 Modifying the ksh History File

The **ksh** stores the commands that you enter at the keyboard in the *.sh_history* file in your home directory by default. You can specify a different history file using the HISTFILE environment variable in your *.profile* file.

For example, to use the file *.history* instead of *.sh_history*, add the following lines to your *.profile*:

```
HISTFILE=~/.history
export HISTFILE
```

You can specify the maximum number of previously entered commands that you can retrieve from the history file with the HISTSIZE environment variable. If HISTSIZE is not set, **ksh** stores 128 commands by default. There is no limit to the number of commands that the **ksh** can store.

Note

If HISTSIZE is very large, **ksh** may be very slow at startup time.

The **ksh** does not delete the history file when you exit; the shell appends and stores commands across login sessions. When you log in, **ksh** deletes any commands in your history file that are older than the last number of commands specified by HISTSIZE.

B.7 Manipulating Commands Wider Than the Screen

The **ksh** allows you to enter commands of up to 256 characters from the terminal. You can define the maximum width of the command-line display (80 columns by default) using the COLUMNS variable. (See the section "Customizing the ksh Environment" in this appendix for more information about setting environment variables.)

If you edit a command that is wider than the command line, **ksh** automatically scrolls the command line horizontally to the left or right of your screen. In the last column on the right side of the screen, **ksh** displays one of the following characters to show that the line is scrolling:

- `<` scrolls to the right (text to the left is not displayed)
- `>` scrolls to the left (text to the right is not displayed)
- `*` text both to the right and to the left is not displayed

For example, if COLUMNS is not set (the width of the command-line display is the default of 80 columns) and your command line is greater than 78 characters wide, **ksh** scrolls the command line left to display the end of the line. To the right of the command line, **ksh** displays the > character.

Horizontal scrolling does not work when a command is being entered initially.

B.8 Using Expanded cd Capabilities

The **ksh** includes expanded functionality for the **cd**(C) command. You can instruct **ksh** to search through a specified list of directories when you enter pathnames that do not begin with the slash (/) character. To do this, set the CDPATH variable in your *.profile* file. (See the section "Customizing the ksh Environment" for more information about setting environment variables.)

The **ksh** provides an option to **cd** that allows you to return quickly to your previous working directory. For example, if you are in the */usr/spool/mail* directory and you enter:

cd /usr/bin

you can return to your previous directory, */usr/spool/mail*, by entering **cd -**. From this directory, you can enter **cd -** again to return to the */usr/bin* directory.

The **ksh** provides a means for changing to a directory with a pathname that is slightly different from your current working directory. To do this, use the following format:

cd *old new*

where *old* is the part of the pathname that you want to change and *new* is what you want to change it to. For example, if you are in */usr/spool/mail* and you want to change to */usr/bin/mail*, enter:

cd spool bin